

Fault Tolerant BFS Structures: A Reinforcement-Backup Tradeoff

Merav Parter ^{*†}

David Peleg ^{*}

April 17, 2015

Abstract

This paper initiates the study of fault resilient network structures that mix two orthogonal protection mechanisms: (a) *backup*, namely, augmenting the structure with many (redundant) low-cost but fault-prone components, and (b) *reinforcement*, namely, acquiring high-cost but fault-resistant components. To study the trade-off between these two mechanisms in a concrete setting, we address the problem of designing a (b, r) *fault-tolerant* BFS (or (b, r) **FT-BFS** for short) structure, namely, a subgraph H of the network G consisting of two types of edges: a set $E' \subseteq E$ of $r(n)$ fault-resistant *reinforcement* edges, which are assumed to never fail, and a (larger) set $E(H) \setminus E'$ of $b(n)$ fault-prone *backup* edges, such that subsequent to the failure of a single fault-prone backup edge $e \in E \setminus E'$, the surviving part of H still contains an BFS spanning tree for (the surviving part of) G , satisfying $\text{dist}(s, v, H \setminus \{e\}) \leq \text{dist}(s, v, G \setminus \{e\})$ for every $v \in V$ and $e \in E \setminus E'$. We establish the following tradeoff between $b(n)$ and $r(n)$: For every real $\epsilon \in [0, 1]$, if $r(n) = \tilde{\Theta}(n^{1-\epsilon})$, then $b(n) = \tilde{\Theta}(n^{1+\epsilon})$ is necessary and sufficient. More specifically, it was shown in [14] that for $\epsilon = 1$, **FT-BFS** structures (with no reinforced edges) require $\Theta(n^{3/2})$ edges, and this number of edges is sufficient. At the other extreme, if $\epsilon = 0$, then $n - 1$ reinforced edges are sufficient with no need for backup. Here, we present a polynomial time algorithm that given an undirected graph $G = (V, E)$, a source vertex s and a real $\epsilon \in [0, 1]$, constructs a $(b(n), r(n))$ **FT-BFS** with $r(n) = O(n^{1-\epsilon})$ and $b(n) = O(\min\{1/\epsilon \cdot n^{1+\epsilon} \cdot \log n, n^{3/2}\})$. We complement this result by providing a nearly matching lower bound, showing that there are n -vertex graphs for which any $(b(n), r(n))$ **FT-BFS** structure requires $\Omega(\min\{n^{1+\epsilon}, n^{3/2}\})$ backup edges when $r(n) = \Omega(n^{1-\epsilon})$ edges are reinforced.

^{*}Department of Computer Science and Applied Mathematics. The Weizmann Institute of Science, Rehovot, Israel. E-mail: {merav.parter,david.peleg}@weizmann.ac.il. Supported in part by the Israel Science Foundation (grant 894/09), and the I-CORE program of the Israel PBC and ISF (grant 4/11).

[†]Recipient of the Google European Fellowship in distributed computing; research is supported in part by this Fellowship.

1 Introduction

Background and Motivation. Modern day communication networks support a variety of logical structures and services, and depend on their uninterrupted operation. Following the immense recent advances in telecommunication networks, the explosive growth of the Internet, and our increased dependence on these infrastructures, guaranteeing the survivability of communication networks has become a major objective in both practice and theory. An important aspect of this objective is *survivable network design*, namely, the design of low cost high resilience networks that satisfy certain desirable *performance requirements* concerning, e.g., their connectivity, distance or capacity. Our focus here, however, is not on planning survivable networks “from scratch”, but rather on settings where an initially existing infrastructure needs to be improved and optimized.


Our interest in this paper is in exploring a natural “quality vs. quantity” tradeoff in survivable network design. Designers and manufactures often face the following design choice when dealing with ensuring product reliability. One option is to invest heavily in the quality and resilience of the various components of the product, making them essentially failure-free. An alternative option is to use unreliable but cheap components, and ensure the reliability of the whole product by employing redundancy, namely, including several “copies” of each component in the design, so that the failure of one component will not disable the operation.

In the context of survivable network design, where the goal is to overcome link disconnections, the “quantity-based” approach to survivability relies on adding to the network many inexpensive (but failure-prone) *backup* links, counting on redundancy to provide resilience and guarantee the desired performance requirements in the presence of failures. In contrast, a “quality-based” approach may rely on *reinforcing* some of the network links, and thus making them failure-resistant (but expensive), counting on these links to ensure the performance requirements. Clearly, these two approaches address two different and *orthogonal* factors affecting the survivability of a network: the *topology*, e.g., the presence of redundant alternate paths, and the *reliability* of individual network components. We would like to study the tradeoff between these two factors in various survivable network design problems.

Towards exploring this tradeoff, we consider the following “mixed” model. Assume that the existing infrastructure consists of a given fixed set V of vertices and a collection E of existing links, and it is required to decide, for each link, among the following three choices: (a) discard the link (in which case it will cost us nothing), (b) purchase it as is (at some low cost B), or (c) “reinforce” it (at some high cost R), making it failure-resilient. The existing initial graph $G(V, E)$ provides a baseline for comparison, in the sense that if we decide on the conservative approach of making no changes, namely, purchasing all the links of the existing network G “as is” (at a cost of $B \cdot |E|$), then the performance properties that can be guaranteed in the presence of failures are those of the existing G . An alternative baseline is obtained by the opposite extreme, namely, basing the design on selecting the smallest subgraph H of G that satisfies the desired performance requirements in the absence of failures, and reinforcing all its links, thus ensuring this performance level.

Unfortunately, both of these two extremes might be too costly. Hence, constructing a survivable subnetwork with a *limited* budget introduces a tradeoff between backup and reinforce-

ment and the system designer is faced with a choice: reinforcing just a few of the links may potentially lead to considerable savings, by allowing one to discard many of the ordinary backup edges and still obtain the same performance properties. To illustrate this point, consider for example an n -vertex network consisting of a single vertex s connected via a single edge e to an $n-1$ -vertex clique (see figure). The edge connectivity of this network is 1, as the removal of e disconnects the graph. Hence the conservative approach of keeping all existing edges leaves this network with a low level of survivability. In contrast, in a mixed model allowing also reinforcements, it is sufficient to reinforce a single edge, namely, e , in order to obtain a high level of survivability, even by purchasing only a fraction of the edges of the clique.



Our Contributions. To initiate the study of the tradeoff between reinforcement and backup in survivable network design, we consider in this paper the concrete problem of designing (in the mixed model) a *fault-tolerant Breadth-First structure* (or **FT-BFS** for short), namely, a subnetwork that preserves distances with respect to a given source vertex s in the presence of an edge failure. Formally, given a network $G(V, E)$ and a source vertex s in G , a $(b(n), r(n))$ -**FT-BFS** is a subgraph H of G consisting of two types of edges: a set $E' \subseteq E$ of $r(n)$ fault-resistant *reinforcement* edges, which are assumed to never fail, and a (larger) set $E(H) \setminus E'$ of $b(n)$ fault-prone *backup* edges, such that subsequent to the failure of a single fault-prone backup edge $e \in E \setminus E'$, the surviving part of H still contains an BFS spanning tree for (the surviving part of) G , satisfying $\text{dist}(s, v, H \setminus \{e\}) \leq \text{dist}(s, v, G \setminus \{e\})$ for every $v \in V$ and $e \in E \setminus E'$. We establish the following tradeoff between $b(n)$ and $r(n)$: For every real $\epsilon \in [0, 1]$, if $r(n) = \tilde{\Theta}(n^{1-\epsilon})$, then $b(n) = \tilde{\Theta}(\min\{n^{1+\epsilon}, n^{3/2}\})$ is necessary and sufficient.

It was shown in [14], that for $\epsilon = 1$, **FT-BFS** structure requires $\Theta(n^{3/2})$ edges. In the other extreme case of $\epsilon = 0$, by reinforcing all the edges of the BFS tree, no backup is needed. This result can also be interpreted in the following manner. Let B (resp., R) be the cost of a backup (resp., reinforced) edge. The total cost of a (b, r) **FT-BFS** structure is given by $B \cdot b(n) + R \cdot r(n) = \tilde{O}(n^{1-\epsilon} \cdot B + n^{1+\epsilon} \cdot R)$. Hence, the minimum cost (b, r) **FT-BFS** obtained by taking $\epsilon = \tilde{O}(\log(R/B)/\log n)$.

We complement the upper bound construction of (r, b) **FT-BFS** structures by presenting a nearly matching lower bound. We show that there are n -vertex graphs for which any $(b(n), r(n))$ **FT-BFS** structure for $r(n) = \Omega(n^{1-\epsilon})$ requires $\Omega(\min\{n^{1+\epsilon}, n^{3/2}\})$ backup edges. In our lower bound constructions, we also consider a generalized structure referred to as a (b, r) *fault-tolerant multi-source BFS tree*, or **FT-MBFS tree** for short, aiming to provide a (b, r) **FT-BFS** structure at each source vertex $s \in S$ for some subset of sources $S \subseteq V$. We show that a (b, r) **FT-MBFS** structure for $r(n) = \Omega(|S|^\epsilon \cdot n^{1-\epsilon})$ requires $b(n) = \Omega(\min\{\sqrt{|S|} \cdot n^{3/2}, |S|^{1-\epsilon} \cdot n^{1+\epsilon}\})$ edges for every $\epsilon \in (0, 1]$.

Techniques and proof outline. Studying (b, r) **FT-BFS** structures significantly differs from their standard **FT-BFS** counterparts (for $r(n) = 0$) in both the upper and lower bounds. Let $\pi(s, v)$ be an $s - v$ shortest-path in G . The initial structure consists of the BFS tree

$T_0 = \bigcup_{v \in V} \pi(s, v)$. It is then augmented by adding to it the last edges of some carefully chosen replacement-paths. For an edge $e \in \pi(s, v)$, a replacement path $P_{v,e}$ is *new-ending* path if its last edge was not present in the structure when the path was selected by the algorithm. A new-ending replacement path $P_{v,e}$ has the following structure. It consists of a prefix of $\pi(s, v)$ followed by a detour D avoiding the failing edge e and joining the $\pi(s, v)$ path at the terminal v . An essential component in our analysis deals with the detour segment of the single failure replacement paths. The analysis of FT-BFS structure [14] focused on a *single* terminal v and showed that it has $O(\sqrt{n})$ new-ending replacement-paths (with distinct last edges). The current setting of (b, r) FT-BFS structures is more involved and requires studying the interactions between detours of *different* vertices. In particular, the current construction has two simultaneous objectives: minimizing the number of backup edges in the structure as well as selecting at most $r(n)$ reinforced edges. In other words, when constructing a (b, r) FT-BFS structure with $o(n^{3/2})$ edges, one has the privilege of discarding the protection against the failure of $r(n)$ edges, which are reinforced.

The upper bound of [14] was achieved by analyzing the interactions between the detours of $s - v$ new-ending replacement-paths P_{v,e_i} and P_{v,e_j} for some $e_i, e_j \in \pi(s, v)$. It was shown that upon a proper construction of the replacement-paths, these detours are vertex disjoint¹, except for the common endpoint v , and hence these detours are vertex-consuming, which enables bounding their number. In contrast, studying (b, r) structures requires understanding the interaction between detours of *distinct* terminals. These detours may overlap and are not necessarily vertex disjoint, hence bounding their number calls for new techniques.

Our key observation is that the interactions (referred hereafter as *interference*) between detours can be roughly classified into two types depending on the relation between the edges protected by the corresponding detours. Each of these interference types gives raise to unique structural characterization and volume constraints that enable us to bound the cardinality of their corresponding paths. The first type of interference concerns $P_{v,e}$ and $P_{t,e'}$ paths both whose failing edges e and e' occur below the least common ancestor of v and t in the BFS tree T_0 . We show that adding the last edges of $O(n^\epsilon)$ such replacement-paths protecting against the failure of the deepest edges of each $s - v$ path is sufficient, i.e., it leaves no unprotected edge (that is protected by a replacement-path of this type) in the structure. We then turn to consider the second type of interference, where at least one of the faulty edges, say e , occurs above the least common ancestor of v and t . Analyzing the interaction between detours that protect edges on the *same* shortest-path turns out to be more involved. Our technique is based on the *heavy-path-decomposition* procedure of Sleator and Tarjan [19] (slightly adapted by Baswana and Khanna [2]), applied on T_0 . This decomposition is obtained by $O(\log n)$ recursive calls on partial trees $T' \subseteq T_0$, where each recursive call results with a collection of paths in T_0 whose edges appear in the $s - v$ shortest-paths of a *distinct* set of vertices. The advantage of this approach is that equipped with our interference classification, the analysis is reduced to solving the subproblem (i.e., designing the (b, r) structure) for the case where the failing events are restricted to a given path $\psi \subseteq T_0$ in the tree-decomposition (a similar approach is taken in [2] for a different problem). In other words, when handling the

¹if their last edge is distinct

second type of interference, there is an *independence* between the tree-decomposition paths $\psi_i, \psi_j \subseteq T_0$ that were generated at the same level of the recursion. Since there are $O(\log n)$ recursion levels, summing over all levels increases our bounds by a logarithmic factor. The final structure H is then given by the union of the substructures for each of the paths in the tree-decomposition². By collecting the last edges of carefully selected replacement-paths protecting the failures on ψ , for every path ψ in the tree-decomposition, it is then shown that there are $\tilde{O}(n^{1-\epsilon})$ unprotected edges in the structure.

Turning to the lower bound, (b, r) FT-BFS structures for large $b(n)$ and $r(n)$ values require a more delicate construction when compared to standard FT-BFS structures. The design of the lower bound graph is governed by two opposing forces whose balance is to be found. Specifically, since detours are vertex consuming, to end up with a dense structure with many backup edges, the detours (and as a result also the shortest-paths) of many vertices should *collide*. For instance, in the lower bound construction of FT-BFS structures, the $s - v$ shortest-path of $\Theta(n)$ vertices is the *same*. In other words, a large number of backup edges implies packing many shortest-paths and detours efficiently. Since the lower bound construction of [14] involved only $\Theta(\sqrt{n})$ edges on the $s - v$ shortest-paths, a new approach is needed when trying to maximize the number of reinforced edges in the structure to $O(n^{1-\epsilon})$ for $\epsilon \in (0, 1/2)$. In particular, large reinforcement forces the construction to distribute the vertices on *distinct* shortest-paths so as to increase the number of edges that have large cost and hence should be reinforced. Our construction then finds the fine balance between these forces, matching our upper bounds up to logarithmic factors.

Related Works. To the best of our knowledge, this paper is the first to study the backup - reinforcement tradeoff in survivable network design for (b, r) FT-BFS structures.

The question of designing sparse FT-BFS structures (without link reinforcement) has been studied in [14], using the notion of *replacement paths*. For a source node s , a target node v and an edge $e \in G$, a *replacement path* is the shortest $s - v$ path $P_{v,e}$ that does not go through e . An FT-BFS structure consists of the collection of all $P_{v,e}$ replacement paths for every target $v \in V$ and edge $e \in E$. It is shown in [14] that for every graph G and source node s there exists a (polynomial time constructible) FT-BFS structure H with $O(n^{3/2})$ edges. This result was complemented by a matching lower bound showing that for every sufficiently large integer n , there exist an n -vertex graph G and a source node $s \in V$, for which every FT-BFS structure is of size $\Omega(n^{3/2})$. Hence the insistence on *exact* distances makes FT-BFS structures significantly denser (hence expensive) compared their fault-prone counterparts (namely, BFS trees). This last observation motivates the idea of studying the mixed model and makes FT-BFS structures an attractive platform for studying the backup-reinforcement tradeoff.

The notion of FT-BFS trees is also closely related to the *single-source replacement paths* problem, studied in [9]. That problem requires to compute the collection $\mathcal{P}(s)$ of all $s - t$ replacement paths $P_{t,e}$ for every $t \in V$ and every failed edge e that appears on the $s - t$ shortest-path in G . The vast literature on *replacement paths* (cf. [3, 9, 17, 20, 21]) focuses

²The paths of the tree-decomposition do not cover all the edges in the BFS tree, however, the remaining uncovered edges can be handled directly by adding the last edges of the corresponding replacement-paths.

on *time-efficient* computation of these paths as well as on their efficient maintenance in data structures (a.k.a *distance oracles*).

Constructions of sparse fault tolerant *spanners* for \mathbb{R}^d Euclidean space were studied in [7, 12, 13]. Algorithms for constructing sparse edge and vertex fault tolerant spanners for arbitrary undirected weighted graphs were presented in [5, 8]. Note, however, that the use of costly link reinforcements for attaining fault-tolerance in spanners is less attractive than for FT-BFS structures, since the cost of adding fault-tolerance via backup edges (in the relevant complexity measure) is often low (e.g., merely polylogarithmic in the graph size n), hence the gains expected from using reinforcement are relatively small.

Constructions of edge fault-tolerant spanners with *additive* stretch are given in [4], and the case of single vertex fault has been recently studied in [16].

Discussion. The presented mixed model implicitly reflects the notion of *economy of scale*, stating that the cost per unit decreases with the amount of purchased units. Indeed, economy of scale has been incorporated explicitly in many network design tasks by using a *sublinear* cost function. These scenarios are known in the literature as *buy-at-bulk* problems [1, 18, 6] in which capacity is sold with a “volume discount”: the more capacity is bought, the cheaper is the price per unit of bandwidth.

Economy of scale arises implicitly in (b, r) FT-BFS structures. The main objective in constructing (b, r) FT-BFS is to minimize the number of backup edges $b(n)$ subject to a bound limit on the number of reinforced edges $r(n) = O(n^{1-\epsilon})$. We say that a vertex $v \in V$ *uses* an edge if the edge is on the $s - v$ shortest path in G (for the sake of discussion, assume that all shortest paths in G are unique). The *cost* of an edge e , $\text{Cost}(e)$, is the number of backup edges required to be added to the structure upon its failing (to provide alternative shortest-paths in the surviving network). Since reinforcement is expensive, it is beneficial to reinforce an edge that has many users (i.e., appears on many $s - v$ shortest paths). The intuition behind this observation is that an edge e with many users may cause a considerable damage upon failing, hence its cost $\text{Cost}(e)$ should be large. Whereas in the traditional backup mechanism $\text{Cost}(e)$ scales with the number of users, in the reinforcement mechanism the reinforcement cost of each edge is fixed and independent of the number of users. Therefore, for certain B/R ratios, the cost of reinforcement becomes sublinear with the number of users. A similar intuition arises in the *single sink rent-or-buy* problem [11].

The problem considered here, namely, the construction of (r, b) FT-BFS structure, deviates from the traditional rent-or-buy and buy-at-bulk network design tasks in the sense that it concerns a fault resilient distance preserving (and not merely connectivity preserving) structure. In addition, whereas most network design tasks are stated as combinatorial optimization problems and their solution employs polyhedral combinatorics, the goal of the current paper is to establish a *universal* bound on the tradeoff between backup and reinforcement.

This paper aims at establishing universal lower and upper bounds for (b, r) FT-BFS structures. In particular, although the universal upper bound is nearly tight (upto logarithmic factors), our upper bound constructions might be far from optimal in some instances (see the example of Fig. 5 in [15]). This motivates the study of (b, r) FT-BFS structures from the

combinatorial optimization point of view. Specifically, two natural optimization problems can be defined within this context. The first (resp., second) formulation aims at optimizing the number of backup edges $b(n)$ (resp., reinforced edges $r(n)$) subject to a given bound on the number of reinforced edges (resp., backup edges). For example, when $r(n) = \tilde{O}(n^{1-\epsilon})$, we know that for every input graph G and source vertex $s \in V$, one can construct a structure with $b(n) = \tilde{O}(n^{1+\epsilon})$ backup edges, and there are graphs for which $b(n) = \Omega(n^{1+\epsilon})$ backup edges are essential. Yet, there are graphs for which $b(n) = O(n)$ backup edges are sufficient and using the current construction is too wasteful.

Aside from optimization tasks for (r, b) FT-BFS structures, the presented reinforcement-backup tradeoff can be studied in a more generalized setting. In fact, it can be integrated into a large collection of survivability network design tasks. We hope that this work will pave the way for studying this setting, leading to new theoretical tools and techniques as well as to a better understanding of fault resilient structures.

2 Preliminaries and Notation

Let $E(v, G) = \{(u, v) \in E(G) \mid u \in V\}$ be the set of edges incident to the vertex v in the graph G and let $\deg(v, G) = |E(v, G)|$ denote the degree of v in G . When the graph G is clear from the context, we may simply write $\deg(v)$ and $E(v)$. For a subgraph $G' = (V', E') \subseteq G$ (where $V' \subseteq V$ and $E' \subseteq E$) and a pair of vertices $u, v \in V$, let $\text{dist}(u, v, G')$ denote the shortest-path distance in edges between u and v in G' . For a path $P = [u_1, \dots, u_k]$, let $\text{LastE}(P)$ denote the last edge of P , let $|P|$ denote the length of P in edges, i.e., $k - 1$, and let $P[u_i, u_j]$ be the subpath of P from u_i to u_j . For paths P_1 and P_2 where the last vertex of P_1 equals the first vertex of P_2 , let $P_1 \circ P_2$ denote the path obtained by concatenating P_2 to P_1 . Throughout, the edges of these paths are considered to be directed away from the source s . Given an $s - t$ path P and an edge $e = (u, v) \in P$, let $\text{dist}(s, e, P)$ be the distance (in edges) between s and e on P . For an edge $e = (u, v) \in T_0$, define $\text{dist}(s, e) = i$ if $\text{dist}(s, u, G) = i - 1$ and $\text{dist}(s, v, G) = i$. For a subset $V' \subseteq V$, let $G(V')$ be the subgraph of G induced by V' . Let $\text{LCA}(u, v)$ be the least common ancestor of u and v in T_0 .

For vertices $u, v \in V$ and subgraph $G' \subseteq G$, let $SP(u, v, G')$ be the collection of all $s - v$ shortest-path in G' , i.e., $|P| = \text{dist}(s, v, G')$ for every $P \in SP(s, v, G')$. For a positive weight assignment $W : E(G) \rightarrow \mathbb{R}_{>0}$, let $SP(s, v, G', W)$ be the collection of $s - v$ shortest-paths in G' according to the weights of W . In this paper, the weight assignment W is chosen as to guarantee the uniqueness the shortest-paths in every $G' \subseteq G$. That is W is used to break to shortest-path ties in G' in a consistent manner. In such a case, we override notation and let $SP(s, v, G', W) \in SP(s, v, G')$ be the unique $s - v$ shortest-path in G' with the weights of W . Given a source vertex s and target vertex v , let $\pi(s, v) = SP(s, v, G, W)$ be the unique $s - v$ path in G according to W . Define $T_0(s) = \bigcup_{v \in V} \pi(s, v)$ as the BFS tree rooted at s . When the source s is clear from the context, we simply write T_0 .

For a vertex v and an edge e , each path in $SP(s, v, G \setminus \{e\})$ is referred to as a *replacement-path*. Note that if $e \notin \pi(s, v)$, then $\pi(s, v)$ is a replacement path as it appears in $SP(s, v, G \setminus \{e\})$. A vertex w is a *divergence point* of the $s - v$ paths P_1 and P_2 if $w \in P_1 \cap P_2$ but the

next vertex u after w (i.e., such that u is closer to v) in the path P_1 is not in P_2 .

ϵ **FT-BFS** and *protected edges*. For a subgraph $H \subseteq G$ and a source vertex s , an edge e is *protected* in H if $\text{dist}(s, v, H \setminus \{e\}) = \text{dist}(s, v, G \setminus \{e\})$ for every $v \in V$ and otherwise it is *unprotected*. In other words, the edge e is protected if for every vertex v , H contains at least one replacement path $P_{v,e} \in SP(s, v, G \setminus \{e\})$.

Definition 2.1 (ϵ **FT-BFS**). *For every real $\epsilon \in [0, 1]$, a subgraph $H \subseteq G$ is an ϵ **FT-BFS** with respect to s , if it contains $O(n^{1-\epsilon})$ unprotected edges. That is, there exists a subset of $O(n^{1-\epsilon})$ edges E' such that $\text{dist}(s, v, H \setminus \{e\}) = \text{dist}(s, v, G \setminus \{e\})$ for every $v \in V$ and $e \in E(G) \setminus E'$. Alternately, an ϵ **FT-BFS** H can be thought of as a (b, r) **FT-BFS** taking E' to be the set of $r(n) = O(n^{1-\epsilon})$ reinforcement edges and $E(H) \setminus E' = O(b(n))$ to be the backup edges.*

Note that in the context of the reinforcement-backup model, unprotected edges are viewed as edges that should be reinforced in the structure, since by definition, in the (b, r) **FT-BFS** structure, all backup edge are *protected*, and unprotected edges are not allowed to exist.

We now define a more refined notion of protected edges that is determined by the existence of the *last edges* of the replacement-paths in the subgraph H (instead of requiring the existence of the entire replacement path in H). Given a subgraph $H \subseteq G$, we say that the edge e is *v -last-unprotected* in H if there exists no replacement path $P_{v,e} \in SP(s, v, G \setminus \{e\})$ whose last edge $\text{LastE}(P_{v,e})$ is in H , otherwise the edge is *v -last-protected*. An edge e is *last-unprotected* in H , if there exists at least one vertex $v \in V$ for which e is *v -last-unprotected*, otherwise it is *last-protected*. Note that the notion of protected edge refers to the case where every vertex v has at least one $s - v$ replacement-path protecting against the failing of e in H . In contrast, the notion of last-protected edges refers to the existence of the last edge of these replacement-path (and not the entire path) in the subgraph H . The next observation relates the properties of “last-protected” and “protected”.

Observation 2.2. *If e is last-protected in H , then e is protected, i.e., $\text{dist}(s, v, H \setminus \{e\}) = \text{dist}(s, v, G \setminus \{e\})$, $\forall v \in V$.*

Proof: Let e be a last-protected edge in H . Assume towards contradiction that the claim does not hold and let

$$BV = \{v \mid \text{dist}(s, v, H \setminus \{e\}) > \text{dist}(s, v, G \setminus \{e\})\}$$

be the set of “bad vertices,” namely, vertices for which the $s - v$ shortest path distance in $H \setminus \{e\}$ is greater than that in $G \setminus \{e\}$. (By the contradictory assumption, it holds that $BV \neq \emptyset$.) For every vertex v , let $P_{v,e}^* \in SP(s, v, G \setminus \{e\})$ be a replacement-path satisfying that $\text{LastE}(P_{v,e}^*) \in H$, (since e is *v -last-protected* for every v , this path exists). Define $BE(v) = P_{v,e}^* \setminus E(H)$ to be the set of “bad edges,” namely, the set of $P_{v,e}^*$ edges that are missing in H . By definition, $BE(v) \neq \emptyset$ for every bad vertex $v \in BV$. Let $\tilde{d}(v) = \max_{e \in BE(v)} \{\text{dist}(s, e, P_{v,e}^*)\}$ be the maximal depth of a missing edge in $BE(v)$, and let $DM(v)$ denote that “deepest missing edge” for v , i.e., the edge e on $P_{v,e}^*$ satisfying

$\text{dist}(s, e, P_{v,e}^*) = \tilde{d}(v)$. Finally, let $v' \in BV$ be the vertex that minimizes $\tilde{d}(v)$, and let $e_1 = (x, y) \in BE(v')$ be the deepest missing edge on $P_{v',e}^*$, namely, $e_1 = DM(v')$. Note that e_1 is the *shallowest* “deepest missing edge” over all bad vertices $v \in BV$. Let $P_1 = P_{s,y,e}^*$, $P_2 = P_{s,y,e}^*[s, y]$ and $P_3 = P_{s,v',e}^*[y, v']$. Note that since $v' \in BV$, it follows that also $y \in BV$. (Otherwise, if $y \notin BV$, then any s – y shortest-path $P' \in SP(s, y, H \setminus \{e\})$, where $|P'| = |P_{y,e}^*|$, can be appended to P_3 resulting in $P'' = P' \circ P_3$ such that (1) $P'' \subseteq H \setminus \{e\}$ and (2) $|P''| = |P'| + |P_3| = |P_2| + |P_3| = |P_{s,v',e}^*|$, contradicting the fact that $v' \in BV$.) We conclude that $y \in BV$. Finally, note that $\text{LastE}(P_1) \in H$ by definition, and therefore the deepest missing edge of y must be shallower, i.e., $\tilde{d}(y) < \tilde{d}(v')$. However, this is in contradiction to our choice of the vertex v' . The lemma follows. \blacksquare

3 Algorithm

In this section, we describe a construction of an ϵ FT-BFS subgraph H containing $O(1/\epsilon' \cdot \log n \cdot n^{1+\epsilon'})$ edges where $\epsilon' = \epsilon + \log(\log n / \epsilon) / \log n$ for every $\epsilon \in (0, 1]$. In the next section we prove the following.

Theorem 3.1. *For every input n -vertex graph G , source vertex s and real $\epsilon \in (0, 1]$, there exists an ϵ FT-BFS $H \subseteq G$ with $O(\min\{1/\epsilon' \cdot \log n \cdot n^{1+\epsilon'}, n^{3/2}\})$ edges. Hence, in particular, one can construct a (b, r) FT-BFS structure with $b(n) = O(\min\{1/\epsilon \cdot n^{1+\epsilon} \cdot \log n, n^{3/2}\})$ and $r(n) = O(1/\epsilon \cdot n^{1-\epsilon} \cdot \log n)$.*


By [14], there exists a polynomial time algorithm for constructing FT-BFS structures with $O(n^{3/2})$ edges, hence the claim holds trivially for $\epsilon \geq 1/2$, and to establish the theorem, it remains to consider the case where $\epsilon \in (0, 1/2)$. To deal with this case, we next describe an explicit construction for an ϵ FT-BFS H which is then analyzed in the following subsection.

3.1 Phase (S0): Preprocessing

Algorithm Pcons for constructing the replacement-paths. The goal of the preprocessing phase (S0) is to define a function $\mathcal{RP} : (V \times E) \rightarrow E$ that maps each vertex-edge pair $\langle v, e \rangle$ to a replacement path $P_{v,e} \subseteq E$. These paths will be used in the main construction. Let T_0 be a BFS tree rooted at s in G . Algorithm Pcons iterates over every vertex $v \in V$ and every edge $e \in \pi(s, v)$. For a given pair $\langle v, e \rangle$, the algorithm first tests if there exists an s – v replacement path whose last edge is already in T_0 . Let $G'(v) = (G \setminus E(v, G)) \cup E(v, T_0)$. Now, if $\text{dist}(s, v, G'(v) \setminus \{e\}) = \text{dist}(s, v, G \setminus \{e\})$, then let $\mathcal{RP}(\langle v, e \rangle) = P_{v,e} = SP(s, v, G'(v) \setminus \{e\}, W)$. Else, (i.e., the replacement-path $P_{v,e}$ must include a new last edge that is not in T_0), the algorithm attempts to select the s – v replacement-path whose divergence point from $\pi(s, v)$ is as close to s as possible. Specifically, let $\pi(s, v) = [u_0 = s, u_1, \dots, u_k = v]$ and $e = (u_i, u_{i+1})$. For every $j \in \{0, \dots, i\}$, define $G_j(v) = G \setminus V(\pi(u_j, u_k)) \cup \{u_j, u_k\}$. Note that $e \notin G_j(v)$. Define j^* as the minimal index j satisfying that $\text{dist}(s, v, G_j(v)) = \text{dist}(s, v, G \setminus \{e\})$ and let $P_{v,e} = SP(s, v, G_{j^*}(v) \setminus \{e\}, W)$.

Replacement-path classification. A replacement-path $P = P_{v,e}$ is *new-ending* if its last edge is not in T_0 . A vertex-edge pair $\langle v, e \rangle$ is *uncovered* if its replacement path $P_{v,e}$ is new-ending. The following observation follows immediately by the construction of Alg. Pcons.

Observation 3.2. Consider a new-ending path $P = P_{v,e}$. Let $d(P)$ be the first divergence point of P and $\pi(s, v)$. Then P can be decomposed into $P = \pi(s, d(P)) \circ D(P)$ where $D(P) = P \setminus E(\pi(s, v))$, referred to as the detour segment, departs from $\pi(s, v)$ at $d(P)$ and returns only at v , i.e., $D(P) = P[d(P), v]$ and $\pi(s, v)$ are vertex disjoint besides the common endpoints $d(P)$ and v .



Let $\mathcal{UP} = \{\langle v, e \rangle \mid \text{LastE}(P_{v,e}) \notin T_0\}$ be the collection of all uncovered vertex-edge pairs. Let $\mathcal{UP}(v) = \{\langle w, e \rangle \in \mathcal{UP} \mid w = v\}$ be the uncovered pairs of v (hence, $\mathcal{UP} = \bigcup_{v \in V} \mathcal{UP}(v)$).

Throughout, we consider the edges of T_0 to be directed away from s , hence referring to the edge $e = (x, y) \in T_0$ implies that $\text{dist}(s, x, G) < \text{dist}(s, y, G)$.


The following definitions are key to in our construction and the subsequent analysis. For two tree edges $e = (a, b), e' = (c, d) \in T_0$, we say that $e \sim e'$ if $\text{LCA}(b, d) \in \{b, d\}$, i.e., $e, e' \in \pi(s, v)$ for $v = \{b, d\} \setminus \{\text{LCA}(b, d)\}$, otherwise $e \not\sim e'$. (In the figure, $e \sim e'$ and $e' \not\sim e''$.) In our construction, we may impose an ordering on a subset of v 's uncovered pairs. For a given subset of v 's pairs $\text{Add}(\mathcal{P}, v) = \{\langle v, e_1 \rangle, \dots, \langle v, e_k \rangle\} \subseteq \mathcal{UP}(v)$, let $\vec{\mathcal{P}}(v)$ be ordered in increasing distance of e_i from v , i.e., $\vec{\mathcal{P}}(v) = \{\langle v, e_{i_1} \rangle, \dots, \langle v, e_{i_k} \rangle\}$ where $\text{dist}(v, e_{i_1}, \pi(s, v)) < \dots < \text{dist}(v, e_{i_k}, \pi(s, v))$.

(\sim)-interference and ($\not\sim$)-interference. The paths $P = P_{v,e}, P' = P_{t,e'}$ for $\langle v, e \rangle, \langle t, e' \rangle \in \mathcal{UP}$ and $v \neq t$ interfere with each other if their detours intersect at some vertex z internal to both, i.e.,

$$V(D(P)) \cap V(D(P')) \not\subseteq \{d(P), d(P'), v, t\}. \quad (1)$$

Note that according to this definition, interference is *symmetric*, i.e., if P interferes with P' then P' interferes with P as well. For every uncovered pair $\langle v, e \rangle \in \mathcal{UP}$, denote the set of pairs $\langle t, e' \rangle$ whose corresponding path $P' = P_{t,e'}$ interferes with $P = P_{v,e}$ by

$$\mathcal{I}(\langle v, e \rangle) = \{\langle t, e' \rangle \in \mathcal{UP} \mid t \neq v, P_{t,e'} \text{ and } P_{v,e} \text{ satisfy Eq. (1)}\}.$$



Our construction is heavily based on distinguishing between two types of interference, depending on the relation of the two failing edges protected by the interfered paths. In particular, if the interfering paths $P_{v,e}$ and $P_{t,e'}$ satisfy that $e \not\sim e'$, then we call it ($\not\sim$)-interference, and if $e \sim e'$, then it is (\sim)-interference. For an illustration, see Fig. 1. Let $\mathcal{I}^\sim(\langle v, e \rangle) = \{\langle t, e' \rangle \in \mathcal{I}(P) \mid e \not\sim e'\}$ be the set of pairs whose corresponding paths




Figure 1: Illustration of the two types of interference. The replacement paths $P_{v,e}$ (solid) and $P_{t,e'}$ are (a) $(\not\sim)$ interfering as $e \not\sim e'$, and (b) (\sim) interfering as $e \sim e'$.

$(\not\sim)$ -interfere with $P_{v,e}$.

A given subset of uncovered pairs $\mathcal{P}' \subseteq \mathcal{UP}$ is called a (\sim) -set if $\mathcal{I}^{\not\sim}(\langle v, e \rangle) \cap \mathcal{P}' = \emptyset$ for every $\langle v, e \rangle \in \mathcal{P}'$. Otherwise, it is called $(\not\sim)$ -set. In other words, in a (\sim) -set there is no $(\not\sim)$ -interference between any pair of paths.

3.2 The main construction

Let us start with an overview of the main construction phases. The initial structure H contains T_0 . In phases (S1) and (S2), we add backup edges to H corresponding to last edge of the new ending replacement paths $P_{v,e}$ so that eventually the set of T_0 edges unprotected by H is bounded by $O(1/\epsilon \cdot n^{1-\epsilon} \cdot \log n)$. (These edges will have to be reinforced; all other edges of H will be taken as backup edges.) The high level idea of our main construction is as follows. First, we divide the uncovered pairs \mathcal{UP} into two sub-sets $(\not\sim)$ -set \mathcal{I}_1 and a (\sim) -set \mathcal{I}_2 , by letting

$$\mathcal{I}_1 = \{\langle v, e \rangle \in \mathcal{UP} \mid \mathcal{I}^{\not\sim}(\langle v, e \rangle) \neq \emptyset\} \text{ and } \mathcal{I}_2 = \mathcal{UP} \setminus \mathcal{I}_1.$$

Phase (S1) starts by setting the first (\sim) -set to be $\mathcal{P}_0^C = \mathcal{I}_2$. Then, Phase (S1) employs an iterative process of $K_\epsilon = O(1/\epsilon)$ iterations. Each of these iterations does the following. For every vertex $v \in V$, the algorithm repeatedly adds the $\lceil n^{1/\epsilon} \rceil$ distinct last edges of the remaining $s - v$ replacement-paths of the uncovered pairs in \mathcal{I}_1 protecting the deepest edges on $\pi(s, v)$. In addition, each such iteration i may yield an additional (\sim) -set, \mathcal{P}_i^C , which would be handled in Phase (S2). Thus, at the end of Phase (S1), we have at most $O(1/\epsilon)$ such (\sim) -sets \mathcal{P}_i^C that *partially* cover the pairs of \mathcal{I}_1 . The last edges of the replacement-paths of the pairs in \mathcal{I}_1 that are not covered by the (\sim) -sets are added to H . Phase (S2) of the algorithm is then devoted for considering the (\sim) -sets (i.e., \mathcal{I}_2 and the $O(1/\epsilon)$ additional (\sim) -sets that were created in Phase (S1)). For each such (\sim) -set \mathcal{P}' and for every vertex v ,

the algorithm adds a collection of $O(n^\epsilon \cdot \log n)$ backup edges corresponding to the last edges of $s - v$ replacement-paths of the pairs in \mathcal{P}' . The analysis shows that after Phase (S2), for each of the $O(1/\epsilon)$ (\sim) -sets \mathcal{P}_i^C the number of edges protected by replacement paths corresponding to the pairs collection \mathcal{P}_i^C that are still unprotected by H is $O(n^{1-\epsilon} \cdot \log n)$. Hence, overall there are at most $O(1/\epsilon \cdot n^{1-\epsilon} \cdot \log n)$ edges that are still unprotected by H . Those edges will have to be reinforced.

We now describe the algorithm in detail.

Phase (S1): Handling the (\sim) -set \mathcal{I}_1 .

The next definition is important in this context. For $v \neq t$, a path $P_{v,e}$ π -intersects with path $P_{t,e'} \in \mathcal{I}^\sim(P_{v,e})$ if the detour of $P_{v,e}$ intersects at least one of the vertices of $\pi(\text{LCA}(v, t), t) \setminus \{\text{LCA}(t, v)\}$, see Fig. 2. Note that this property may *not* be symmetric




Figure 2: Illustration of π -intersection. The replacement-path $P_{v,e}$ (solid) π -intersects $P_{t,e'} \in \mathcal{I}^\sim(P_{v,e})$ (dashed).

(unlike interference). That is, it might be the case that $P_{v,e}$ π -intersects $P_{t,e'}$ but not vice-versa.

The replacement paths of the uncovered pairs in some subset $\mathcal{P}_\ell \subseteq \mathcal{I}_1$ can be roughly classified into three types, termed A, B, and C with respect to \mathcal{P}_ℓ . A replacement-path $P_{v,e}$ for $\langle v, e \rangle \in \mathcal{P}_\ell$ is of *type A with respect to \mathcal{P}_ℓ* if it π -intersects at least one path in $\mathcal{I}^\sim(\langle v, e \rangle) \cap \mathcal{P}_\ell$. Let $\mathcal{P}_\ell^A \subseteq \mathcal{P}_\ell$ be the subset of all pairs whose paths is of type A, i.e.,

$$\mathcal{P}_\ell^A = \{ \langle v, e \rangle \in \mathcal{P}_\ell \mid \exists \langle t, e' \rangle \in \mathcal{I}^\sim(\langle v, e \rangle) \cap \mathcal{P}_\ell, P_{v,e} \text{ } \pi\text{-intersects } P_{t,e'} \} . \quad (2)$$

A replacement-path $P_{v,e}$ for $\langle v, e \rangle \in \mathcal{P}_\ell$ is of *type B with respect to \mathcal{P}_ℓ* , if it is not of type A and it (\sim) -interferes with at least one path $P_{t,e'}$ for $\langle t, e' \rangle \in \mathcal{P}_\ell$ that is not of type A as well, i.e., $\langle t, e' \rangle \in \mathcal{P}_\ell \setminus \mathcal{P}_\ell^A$. In such a case, both $\langle v, e \rangle$ and $\langle t, e' \rangle$ are not in \mathcal{P}_ℓ^A , and hence $P_{v,e}$ does not π -intersect $P_{t,e'}$ and vice-versa, implying that $P_{t,e'}$ is of type B as well. Let \mathcal{P}_ℓ^B be

the collection of the pairs whose corresponding path is of type B; formally

$$\mathcal{P}_\ell^B = \{\langle v, e \rangle \in \mathcal{P}_\ell \setminus \mathcal{P}_\ell^A \mid \mathcal{I}^\mathcal{L}(\langle v, e \rangle) \cap (\mathcal{P}_\ell \setminus \mathcal{P}_\ell^A) \neq \emptyset\}. \quad (3)$$

Finally, a replacement-path $P_{v,e} \in \mathcal{P}_\ell$ is of *type C with respect to \mathcal{P}_ℓ* if it is not of type A or B. Note that such a path $P_{v,e}$ satisfies that the intersection $\mathcal{I}^\mathcal{L}(\langle v, e \rangle) \cap (\mathcal{P}_\ell \setminus \mathcal{P}_\ell^A)$ is empty. (This can happen either because $\mathcal{I}^\mathcal{L}(\langle v, e \rangle) \cap \mathcal{P}_\ell = \emptyset$ or because $(\mathcal{I}^\mathcal{L}(\langle v, e \rangle) \cap \mathcal{P}_\ell) \subseteq \mathcal{P}_\ell^A$.) Let $\mathcal{P}_\ell^C = \mathcal{P}_\ell \setminus (\mathcal{P}_\ell^A \cup \mathcal{P}_\ell^B)$ be the set of pairs whose path is of type C. Let

$$K_\epsilon = \lceil 1/\epsilon \rceil + 2. \quad (4)$$

The uncovered pairs of \mathcal{I}_1 are now partitioned into $K_\epsilon + 1$ subsets: K_ϵ (\sim) -sets $\mathcal{P}_1^C, \dots, \mathcal{P}_{K_\epsilon}^C$ and a subset containing all the remaining pairs $\mathcal{I}'_1 = \mathcal{I}_1 \setminus \bigcup_{i=1}^{K_\epsilon} \mathcal{P}_i^C$. Essentially, the subset \mathcal{I}'_1 is “implicit” and is not actually constructed by the algorithm; it consists of all \mathcal{I}_1 pairs $\langle v, e \rangle$ whose last edge of their path $P_{v,e}$ was added to H during one of the K_ϵ iterations of Phase (S1). The analysis shows that the number of distinct last edges of the replacement paths of \mathcal{I}'_1 that were added into H is bounded by $O(1/\epsilon \cdot n^{1+\epsilon})$.

The partition of \mathcal{I}_1 is conducted in K_ϵ iterations. At the end of each iteration, $O(n^{1+\epsilon})$ distinct last edges of the paths that correspond to the first $\langle v, e \rangle$ pairs from \mathcal{I}_1 (the paths protecting the deepest edges on $\pi(s, v)$) are added to H (and intuitively, the pairs of these replacement paths join \mathcal{I}'_1). Initially, let $\mathcal{P}_1 = \mathcal{I}_1$. For every $i = \{1, \dots, K_\epsilon\}$, the next steps are performed:

- Divide \mathcal{P}_i into the subsets $\mathcal{P}_i^A, \mathcal{P}_i^B$ and \mathcal{P}_i^C (according to Eq. (2,3)).
(* Handling the paths of \mathcal{P}_i^C is deferred to Phase (S2). The following steps attempt to handle the paths of $\mathcal{P}_i^A \cup \mathcal{P}_i^B$. *)
- Let $\vec{\mathcal{P}}_i^J(v) = \{\langle v, e_{i_1} \rangle, \dots, \langle v, e_{i_{k_{J_v}}} \rangle\}$ be the ordered $\langle v, e \rangle$ uncovered pairs of v in \mathcal{P}_i^J for every $v \in V$ and $J \in \{A, B\}$ (in increasing distance of the failing edge e_{i_j} from v).
- Add to H , the $\lceil n^\epsilon \rceil$ distinct last edges of the *first* replacement-paths of the pairs in the ordering $\vec{\mathcal{P}}_i^J(v)$.
- Set $\mathcal{P}_{i+1} = \{\langle v, e \rangle \in \mathcal{P}_i^A \cup \mathcal{P}_i^B \mid \text{LastE}(P_{v,e}) \notin H\}$.

This completes Phase (S1). Observe that a pair $\langle v, e \rangle \in \mathcal{P}_i$ that was classified as, say, type A in iteration i , but was not handled (i.e., its last edge was not added to H), joins \mathcal{P}_{i+1} and is re-classified in iteration $i+1$, where it may be classified differently. In particular, if it gets classified into \mathcal{P}_{i+1}^C , then its handling will be deferred to Phase (S2).

Phase (S2): Handling the remaining (\sim) -sets. The input for this step is a collection of (\sim) multi-sets $\mathcal{S} = \{\mathcal{P}_0^C, \mathcal{P}_1^C, \dots, \mathcal{P}_{K_\epsilon}^C\}$.

Preprocessing Sub-Phase (S2.0): Building tree-decomposition for T_0 . As a preprocessing step for handling the (\sim) -sets, the algorithm begins applying to the BFS tree T_0 the *heavy-path-decomposition* technique presented by Sleator and Tarjan [19] and slightly adapted by

Baswana and Khanna [2]. Using this technique, the tree T_0 is broken into vertex disjoint paths $\mathcal{TD} = \{\psi_1, \dots, \psi_t\}$ that satisfy some desired properties; for an illustration see Fig. 3(b).

Fact 3.3. [2] *There exists an $O(n)$ time algorithm for computing a path ψ in T_0 whose removal splits T_0 into a set of disjoint subtrees $T_0(v_1), \dots, T_0(v_j)$ s.t. for every $1 \leq i \leq j$,*

- (1) $|T_0(v_i)| \leq n/2$ and $\psi \cap T_0(v_i) = \emptyset$, and
- (2) $T_0(v_i)$ is connected to ψ through some edge, hereafter denoted $e(\psi, i)$.

The algorithm of Fact 3.3 is applied recursively on T_0 . The output of this recursive procedure is a collection \mathcal{TD} of paths $\psi \subseteq T_0$ plus a set of T_0 edges $e(\psi, i)$ that glue the paths ψ to the tree T_0 . Let $E^+(\mathcal{TD}) = \bigcup_{\psi \in \mathcal{TD}} E(\psi)$ be the set of tree edges occurring on the paths of the decomposition and let $E^-(\mathcal{TD}) = T_0 \setminus E^+(\mathcal{TD})$ be the collection of “glue” edges. In the next Sub-Phase, the algorithm iterates over all the vertices v and add to the structure H a collection of last edges of the replacement-paths protecting against the failing of the glue edges. (In the analysis it is shown that at most $O(n \log n)$ edges are added due to this step.)

Sub-Phase (S2.1): Edge addition based on tree-decomposition [for fixed v]. Define the last edges of the new-ending replacement paths protecting the glue edges $E^-(\mathcal{TD}) \cap \psi(s, v)$ by

$$\widehat{E}(\mathcal{TD}, v) = \{\text{LastE}(P_{v,e}) \mid \langle v, e \rangle \in \mathcal{UP} \text{ and } e \in E^-(\mathcal{TD})\}.$$

Add $\widehat{E}(\mathcal{TD}, v)$ to H .

We now turn to consider the main part of Phase S2. The algorithm treats each (\sim) -set $\mathcal{P} \in \mathcal{S}$ separately, by adding into H $O(n^{1+\epsilon} \cdot \log n)$ distinct last edges of replacement paths carefully selected from the uncovered pairs of \mathcal{P} . In the analysis section, we then show that the total number of edges e with a pair $\langle v, e \rangle$ in \mathcal{P} that are unprotected by H is bounded by $O(n^{1-\epsilon} \cdot \log n)$, and since there are $K_\epsilon + 1 = O(1/\epsilon)$ sets in \mathcal{S} (see Eq. (4)), overall there are $O(1/\epsilon \cdot n^{1-\epsilon} \cdot \log n)$ edges in T_0 that are unprotected by H (and will have to be reinforced).

The selection of the uncovered pairs $\langle v, e \rangle$ whose last edge of their replacement path $P_{v,e}$ is to be added into H is performed in the following manner. The algorithm iterates over every (\sim) -set $\mathcal{P} \in \mathcal{S}$ and every vertex $v \in V$, and selects a subset $\text{Add}(\mathcal{P}, v)$ of v ’s uncovered pairs from \mathcal{P} , where the total number of last edges of their corresponding replacement paths is bounded by $O(n^\epsilon \cdot \log n)$, and then adds these last edges to H . The selection, for \mathcal{P} and v , of pairs to be included in $\text{Add}(\mathcal{P}, v)$ is done in two main phases.

Sub-Phase (S2.2)[for fixed \mathcal{P}, v]: Covering pairs based on shortest-path decomposition into $O(\log n)$ fragments. The $s - v$ shortest-path $\pi(s, v) = [s = u_0, \dots, u_k = v]$ is decomposed into $k' = \lfloor \log |\pi(s, v)| \rfloor$ subsegments of exponentially decreasing length, i.e., where each subsegment consists of the first half of the remaining $\pi(s, v)$ path. For an illustration see Fig. 3(a). Formally, letting u_{i_j} be the vertex at distance $\left\lceil \sum_{\ell=1}^j (|\pi(s, v)|/2^\ell) \right\rceil$ from s on $\pi(s, v)$ for $j \in \{1, \dots, k'\}$, and $u_{i_0} = u_0$, the j ’th subsegment is given by $\pi_j(s, v) = \pi(u_{i_{j-1}}, u_{i_j})$ for

every $j \in \{1, \dots, k'\}$. It then holds that

$$|\pi_j(s, v)| \geq \lfloor |\pi(s, v)| / 2^{j-1} \rfloor \quad \text{and} \quad \sum_{j' > j} |\pi_{j'}(s, v)| \geq |\pi_j(s, v)| / 2 \quad (5)$$

For each of the k' subsegments $\pi_j(s, v)$, let $\mathcal{P}_j(v) = \{\langle v, e \rangle \in \mathcal{P} \mid e \in \pi_j(s, v)\}$ be the set of v 's uncovered pairs from \mathcal{P} whose paths protect the edges in $\pi_j(s, v)$; let $\mathcal{LE}_j(\mathcal{P}, v) = \{\text{LastE}(P_{v,e}) \mid \langle v, e \rangle \in \mathcal{P}_j(v)\}$ be the corresponding last edges of these replacement paths.

A subsegment $\pi_j(s, v)$ is *heavy* with respect to \mathcal{P} if $|\mathcal{LE}_j(\mathcal{P}, v)| \geq \lceil n^\epsilon \rceil$, otherwise it is *light*. For every light subsegment $\pi_j(s, v)$, add $\mathcal{P}_j(v)$ to the collection of selected pairs $\text{Add}(\mathcal{P}, v)$ whose last edges (of their corresponding replacement paths) are later added to H . I.e., add to $\text{Add}(\mathcal{P}, v)$ the pairs $\bigcup \mathcal{P}_j(v)$, where the union is over all $j \in \{1, \dots, k'\}$ s.t. $\pi_j(s, v)$ is light. In addition, we move to $\text{Add}(\mathcal{P}, v)$ some additional pairs $\langle v, e_j^* \rangle$ as follows. For every $j \in \{1, \dots, k'\}$, let e_j^* be the first edge on $\pi_j(s, v)$ (closest to s) such that $\langle v, e_j^* \rangle$ is in $\mathcal{P}_j(v)$. The algorithm adds $\bigcup_{j=1}^{k'} \langle v, e_j^* \rangle$ to $\text{Add}(\mathcal{P}, v)$. (This addition would ensure that the divergence point $d(P_{v,e})$ of the replacement paths $P_{v,e}$ protecting edges e on the segment $\pi_j(s, v)$ and whose last edge was not added to the output structure H , is located inside the segment $\pi_j(s, v)$.)

Sub-Phase (S2.3): Covering pairs depending on both the tree-decomposition and $\pi(s, v)$ decomposition [for fixed \mathcal{P}, ψ, v].

Define $E(\psi, \mathcal{P}, v) = \{e \mid \langle v, e \rangle \in \mathcal{P} \text{ and } e \in \psi \cap \pi(s, v)\}$. Let e^* be the upmost edge in $E(\psi, \mathcal{P}, v)$ (i.e., closest to s). Add $\langle v, e^* \rangle$ to $\text{Add}(\mathcal{P}, v)$.

Next, consider the intersection of ψ with $\pi(s, v)$. Recall that in Sub-Phase (S2.1), the $s - v$ path $\pi(s, v)$ was decomposed into $k' = \lfloor \log |\pi(s, v)| \rfloor$ segments $\pi_1(s, v), \dots, \pi_{k'}(s, v)$. Let $\pi_U(\psi, v)$ be the first, i.e., closest to s , subsegment of $\pi(s, v)$ that intersects ψ such that $\pi_U(\psi, v) \not\subseteq \psi$ and $\pi_U(\psi, v) \cap \psi \neq \emptyset$ (if such exists). Similarly, let $\pi_L(\psi, v)$ be the last, i.e., closest to v , subsegment of $\pi(s, v)$ that intersects ψ such that $\pi_L(\psi, v) \not\subseteq \psi$ and $\pi_L(\psi, v) \cap \psi \neq \emptyset$. See Fig. 4 for an illustration.

Let $\mathcal{P}_U(\psi, v) = \{\langle v, e \rangle \in \mathcal{P} \mid e \in \pi_U(\psi, v) \cap \psi\}$ be the pairs in \mathcal{P} whose replacement paths protect against the failing of the edges in the intersection $\pi_U(s, v) \cap \psi$ and let $\mathcal{LE}_U(\mathcal{P}, \psi, v) = \{\text{LastE}(P_{v,e}) \mid \langle v, e \rangle \in \mathcal{P}_U(\psi, v)\}$ be the last edges of the corresponding replacement paths. If $|\mathcal{LE}_U(\mathcal{P}, \psi, v)| \leq \lceil n^\epsilon \rceil$, then add $\mathcal{P}_U(\psi, v)$ to $\text{Add}(\mathcal{P}, v)$. Finally, let e_U^* be the upmost edge on $\pi(s, v)$ with a pair $\langle v, e_U^* \rangle \in \mathcal{P}_U(\psi, v)$. Then, add $\langle v, e_U^* \rangle$ to $\text{Add}(\mathcal{P}, v)$. The set $\mathcal{P}_L(\psi, v)$ is handled in the same manner as $\mathcal{P}_U(\psi, v)$.

Finally, for every \mathcal{P} and v and for every edge e such that $\langle v, e \rangle \in \text{Add}(\mathcal{P}, v)$ add the last edge of $P_{v,e}$ to H . While the resulting sets $\text{Add}(\mathcal{P}, v)$ might contain many pairs, in the analysis section we show that the total number of new backup edges that will be added to H as a result of these pairs will be at most $O(\log n \cdot n^\epsilon)$ per vertex v and (\sim) -set \mathcal{P} . This completes the description of the algorithm.




Figure 3: Decomposition of shortest-paths and trees. (a) decomposition of the $s - v$ path $\pi(s, v)$ into $O(\log n)$ subsegment. The segment i , $\pi_i(s, v)$ contains roughly half of the remaining vertices. (b) Heavy-tree-decomposition \mathcal{TD} on the BFS tree T_0 according to the algorithm of [2]. The solid edges correspond to the paths of \mathcal{TD} .




Figure 4: Illustration of the intersection between $s_\psi - t_\psi$ path ψ in the tree decomposition \mathcal{TD} and the shortest-path $\pi(s, v)$ decomposition. The vertex w is the LCA of t_ψ and v in T_0 . The intersection of the paths $\psi \cap \pi(s, v)$ is shown in red. The upper (resp., lower) intersected segment is $\pi_U(s, v)$ (resp., $\pi_L(s, v)$).

4 Analysis

4.1 Size Bound

We start with size analysis and use the following fact.

Fact 4.1 ([2]). *For every node $v \in V$,*

(a) $\pi(s, v) \cap E^-(\mathcal{TD}) = O(\log n)$, and (b) $\pi(s, v)$ intersects at most $O(\log n)$ paths in \mathcal{TD} [Lemma 3.6 of [2]].

Lemma 4.2. $|E(H)| = \min\{O(1/\epsilon \cdot n^{1+\epsilon} \cdot \log n), n^{3/2}\}$.

Proof: For $\epsilon \geq 1/2$, the claim trivially holds by [14]. From now on, consider $\epsilon \in (0, 1/2)$. By Fact 4.1(a), the set of edges $\widehat{E}(\mathcal{TD}, v)$ that was added in Sub-Phase (S2.1) contains $O(\log n)$ edges. We now focus on a specific vertex v and (\sim) -set \mathcal{P} and bound the number of new edges corresponding to the pairs of $\text{Add}(\mathcal{P}, v)$ that were collected in Sub-Phase (S2.2-3).

In Sub-Phase (S2.2), the algorithm adds the pairs of the light subsegments $\pi_j(s, v)$. Since there are $O(\log n)$ subsegments and as the number of last edges of replacement paths protecting the edges of a light subsegment is bounded by $O(n^\epsilon)$ edges, overall $O(\log n \cdot n^\epsilon)$ edges are added due to these pairs.

In Sub-Phase (S2.3) we restrict attention to a specific path $\psi \in \mathcal{TD}$ and consider the intersection of $\pi(s, v)$ and ψ . By Fact 4.1(b), every path $\pi(s, v)$ intersects with $O(\log n)$ paths ψ in \mathcal{TD} . Since the algorithm adds the last edges of replacement paths protecting edges on $\pi_U(s, v)$ and $\pi_L(s, v)$ only if their number is bounded by $O(n^\epsilon)$, overall $O(\log n \cdot n^\epsilon)$ edges are added due to this sub-phase. Finally, the total number of pairs $\langle v, e_j^* \rangle$ and $\langle v, e_U^* \rangle, \langle v, e_L^* \rangle$ that are added in Sub-Phase (S2.2-3) is bounded by $O(\log n)$. Altogether, we get that the pairs of $\text{Add}(\mathcal{P}, v)$ contributes $O(\log n \cdot n^\epsilon)$ edges to H . The lemma follows by summing over all n vertices and the $O(1/\epsilon)$ (\sim) sets. ■

We proceed by presenting some useful properties of the paths constructed by Alg. Pcons.

4.2 Basic Replacement Path Properties

Lemma 4.3. *For every $v \in V$, $e \in \pi(s, v)$, it holds that $P_{v,e} \in SP(s, v, G \setminus \{e\})$.*

Proof: If the replacement path $P_{v,e}$ is not new-ending, i.e., $\text{LastE}(P_{v,e}) \in T_0$, then the correctness follows immediately. The interesting remaining case is where the replacement-path $P_{v,e}$ had to include a new-edge that was not in T_0 . We show that in this case, there exists an $s-v$ shortest path in $G \setminus \{e\}$ with a unique divergence point from $\pi(s, v)$ that occurs above the failing edge e . In particular, such a path can be given by letting $P' = SP(s, v, G \setminus \{e\}, W)$. To see this, assume towards contradiction that the divergence point of P' from $\pi(s, v)$ is not unique. Let w_1 (resp., w_2) be the first (resp., second) divergence point of P' from $\pi(s, v)$. There are two cases. If $e \in \pi(w_1, w_2)$, then $P'[w_2, v] = SP(w_2, v, G \setminus \{e\}, W) = \pi(w_2, v)$, in contradiction to the fact that $P_{v,e}$ is new-ending. Otherwise, $e \in \pi(w_2, v)$ and $P'[w_1, w_2] = SP(w_1, w_2, G \setminus \{e\}, W) = \pi(w_1, w_2)$, contradicting the fact that w_1 is a divergence point

from $\pi(s, v)$. This establishes the claim that there exists a replacement-path with a unique divergence point. Since the algorithm picks the replacement-path whose unique divergence point from $\pi(s, v)$ is as close to s as possible, the correctness follows. ■

Recall that for a new-ending path $P_{v,e}$ (i.e., $\langle v, e \rangle \in \mathcal{UP}$), $d(P_{v,e})$ is the first divergence point of $P_{v,e}$ from $\pi(s, v)$. By the construction of the new-ending paths, we have the following.

Claim 4.4. *For every new-ending path $P_{v,e}$:*

- (1) *the divergence point $d(P_{v,e})$ is unique;*
- (2) *there exists no $s-v$ replacement path in $G \setminus \{e\}$ whose unique divergence point is above $d(P_{v,e})$ on $\pi(s, v)$ (i.e., closer to s).*

Proof: Begin with (1). Consider a new-ending path $P_{v,e}$ and let $\pi(s, v) = [u_0 = s, u_1, \dots, u_k = v]$. Recall that for every $j \in \{0, \dots, k-1\}$, $G_j(v) = G \setminus V(\pi(u_j, u_k)) \cup \{u_j, u_k\}$. Let j^* be the minimal index j satisfying that $\text{dist}(s, v, G_j(v) \setminus \{e\}) = \text{dist}(s, v, G \setminus \{e\})$. Alg. Pcons defines $P_{v,e} = SP(s, v, G_{j^*}(v) \setminus \{e\}, W)$. We now show that u_{j^*} is the unique divergence point by showing that $D(P_{v,e}) = P_{v,e}[u_{j^*}, v]$ and $\pi(s, v)$ are vertex disjoint except for the common endpoints u_{j^*} and v . Assume towards contradiction otherwise, and let $u_\ell \in (V(D(P_{v,e})) \cap V(\pi(s, v))) \setminus \{u_{j^*}, v\}$ be the last vertex (closest to v) on $D(P_{v,e})$ that occurs on $\pi(s, v) \setminus \{v\}$. By the definition of $P_{v,e}$, u_ℓ occurs above u_{j^*} on $\pi(s, v)$ and hence also above failing edge e on $\pi(s, v)$. Consider the path $P' = \pi(s, u_\ell) \circ P_{v,e}[u_\ell, v]$. By the selection of u_ℓ , $P' \in SP(s, v, G \setminus \{e\})$ and u_ℓ is the unique divergence point of P' and $\pi(s, v)$. In particular, $P' \subseteq G_\ell$ where $\ell < j^*$. Contradiction to the selection of j^* . Part (1) follows. Part (2) follows immediately by Part (1) and the definition of the paths by Alg. Pcons. ■

Claim 4.5. *Consider two new-ending $s-v$ replacement paths $P_{v,e_{i_1}}, P_{v,e_{i_2}}$ such that $\text{LastE}(P_{v,e_{i_1}}) \neq \text{LastE}(P_{v,e_{i_2}})$ where without loss of generality $e_{i_1} = (x_{i_1}, y_{i_1})$ is above (closer to s) $e_{i_2} = (x_{i_2}, y_{i_2})$ on $\pi(s, v)$. Then, $d(P_{v,e_{i_2}}) \in \pi(y_{i_1}, x_{i_2})$.*

Proof: Towards contradiction, assume otherwise. It then holds that $d(P_{v,e_{i_2}}) \in \pi(s, x_{i_1})$. Since the detour segments are edge disjoint with $\pi(s, v)$ (see Cl. 4.4(1)), we get that there are two $s-v$ paths in $G \setminus \{e_{i_1}, e_{i_2}\}$ given by $P_{v,e_{i_1}}, P_{v,e_{i_2}}$, and the optimality of these paths implies $|P_{v,e_{i_1}}| = |P_{v,e_{i_2}}|$. There are three cases.

Case 1: $d(P_{v,e_{i_1}}) = d(P_{v,e_{i_2}})$. In this case, by the uniqueness of the weight assignment W , we get that $P_{v,e_{i_1}} = P_{v,e_{i_2}}$, in contradiction to the fact that $\text{LastE}(P_{v,e_{i_1}}) \neq \text{LastE}(P_{v,e_{i_2}})$. Case 2: $d(P_{v,e_{i_1}})$ is above $d(P_{v,e_{i_2}})$. In this case, we get a contradiction to Cl. 4.4(2) with respect $P_{v,e_{i_2}}$. Case 3: $d(P_{v,e_{i_2}})$ is above $d(P_{v,e_{i_1}})$. In this case, we get a contradiction to Cl. 4.4(2) with respect $P_{v,e_{i_1}}$. The claim follows. ■

Claim 4.6. *For every $P = P_{v,e}$ such that $\langle v, e \rangle \in \mathcal{UP}(v)$,*

- (1) $|D(P)| = \Omega(\text{dist}(e, v, \pi(s, v)))$.
- (2) *For every $P' = P_{v,e'}$ $\langle v, e' \rangle \in \mathcal{UP}(v)$ satisfying that $\text{LastE}(P) \neq \text{LastE}(P')$ it holds that $D(P') \cap D(P) = \{v\}$.*

Proof: Since $d(P)$ is above the edge e on $\pi(s, v)$, it holds that $|D(P)| = |P(d(P), v)| \geq \text{dist}(e, v, \pi(s, v))$. Consider (2) and assume, towards contradiction, that there exists a mutual vertex $w \in (D(P') \cap D(P)) \setminus \{v\}$. Since $d(P)$ and $d(P')$ are unique divergence points, it holds that $P[w, v] \cap V(\pi(s, v)) = \{v\}$ and $P'[w, v] \cap V(\pi(s, v)) = \{v\}$. Hence,

$$P[w, v] = SP(w, v, (G \setminus V(\pi(s, v))) \cup \{v\}, W) = P'[w, v],$$

in contradiction to the fact that $\text{LastE}(P) \neq \text{LastE}(P')$. ■

4.3 Bounding the number of T_0 edges unprotected by H

Throughout, we consider the final structure H (obtained by the end of Phase (S2)) and denote the path $P_{v,e}$ as H -new-ending if $\text{LastE}(P_{v,e}) \notin H$. Let $\mathcal{UP}(H) = \{\langle v, e \rangle \mid \text{LastE}(P_{v,e}) \notin H\}$ be the uncovered pairs in the final structure H .

For every (\sim) -set $\mathcal{P} \in \mathcal{S}$, let $\mathcal{P}_{\text{miss}} = \{\langle v, e \rangle \in \mathcal{P} \cap \mathcal{UP}(H)\}$ be the pairs of \mathcal{P} that are uncovered by H . Let $E_{\text{miss}}(\mathcal{P}, v) = \{e \mid \langle v, e \rangle \in \mathcal{P}_{\text{miss}}\}$ be the set of edges on $\pi(s, v)$ such that the last edge of the replacement paths of \mathcal{P} pairs were not added to H . Let

$$E_{\text{miss}}(\mathcal{P}) = \bigcup_{v \in V} E_{\text{miss}}(\mathcal{P}, v) \quad (6)$$

be the collection of T_0 edges unprotected by H , corresponding to the paths of \mathcal{P} and let $E_{\text{miss}}(H) = \{e \mid \exists v \text{ s.t. } \langle v, e \rangle \in \mathcal{UP}(H)\}$ be the set of T_0 edges that are unprotected by H . Toward the end of this section, we show that

Lemma 4.7. $|E_{\text{miss}}(H)| = O(1/\epsilon \cdot n^{1-\epsilon} \cdot \log n)$.

The analysis proceeds in two steps. Let $\mathcal{P}^C = \bigcup_{i=1}^{K_\epsilon} \mathcal{P}_i^C$ be the collection of pairs whose corresponding paths are of type C defined in Phase (S1). First, we show that due to Phase (S1), $\mathcal{I}_1 \setminus \mathcal{P}^C$ contains *no* uncovered pair in H , i.e., there is no pair $\langle v, e \rangle \in \mathcal{I}_1 \setminus \mathcal{P}^C$ such that $P_{v,e}$ is H -new-ending path. This implies that it suffices to consider the uncovered pairs of \mathcal{P}_i^C , since $\mathcal{UP}(H) = \bigcup_{\mathcal{P} \in \mathcal{S}} \mathcal{P}_{\text{miss}}$. In the second step, we complete the argument by showing that for each of the $O(1/\epsilon)$ (\sim) -sets \mathcal{P} , the cardinality of $E_{\text{miss}}(\mathcal{P})$, the set of T_0 edges that are unprotected by H , is bounded by $O(n^{1-\epsilon} \cdot \log n)$. Since there are $O(1/\epsilon)$ such sets, overall, we get that

$$|E_{\text{miss}}(H)| = \left| \bigcup_{\mathcal{P} \in \mathcal{S}} E_{\text{miss}}(\mathcal{P}) \right| = O(1/\epsilon \cdot n^{1-\epsilon} \cdot \log n)$$

as desired. We now describe the analysis in detail.

4.3.1 Analysis of Phase (S1)

We begin by establishing a property that holds for every two pairs $\langle v, e_1 \rangle, \langle v, e_2 \rangle \in \mathcal{P}_i^J$ for $J \in \{A, B\}$ such that $\text{LastE}(P_{v,e_1}) \neq \text{LastE}(P_{v,e_2})$. This property plays a key role in our analysis and justifies the classification of the paths of \mathcal{P}_i pairs into the three types.

Lemma 4.8. *Let $P_1 = P_{v,e_1}, P_2 = P_{v,e_2}$ be such that $e_1 = (x_1, y_1)$ is above $e_2 = (x_2, y_2)$ on $\pi(s, v)$, $\text{LastE}(P_1) \neq \text{LastE}(P_2)$, $\langle v, e_2 \rangle \in \mathcal{P}_i^J$ for some $J \in \{A, B\}$ and $i \in \{2, \dots, K_\epsilon\}$. Then there exist a vertex t and an edge $e' \in \pi(s, t)$ satisfying (see Fig. 5)*

(a) $\langle t, e' \rangle \in \mathcal{P}_i$ and hence also $\langle t, e' \rangle \in \mathcal{P}_{i-1}^A \cup \mathcal{P}_{i-1}^B$,

(b) $\text{LCA}(t, v) \in \pi(y_1, x_2)$.

Proof: Let $\mathcal{I}_i(\langle v, e_2 \rangle) = \mathcal{I}^\mathcal{A}(\langle v, e_2 \rangle) \cap \mathcal{P}_i$. Since $\langle v, e_2 \rangle \in \mathcal{P}_i^A \cup \mathcal{P}_i^B$, by Eq. (2) and (3), we get that $\mathcal{I}_i(\langle v, e_2 \rangle) \neq \emptyset$ and $\mathcal{I}_i(\langle v, e_2 \rangle) \subseteq \mathcal{P}_{i-1}^A \cup \mathcal{P}_{i-1}^B$.

To identify the path $P' = P_{t,e'}$ where $\langle t, e' \rangle \in \mathcal{I}_i(\langle v, e_2 \rangle)$, consider two cases depending on the type of the path P_2 with respect to \mathcal{P}_i . *Case 1:* $\langle v, e_2 \rangle \in \mathcal{P}_i^A$ (i.e., P_2 is of type A). Let $\langle t, e' \rangle \in \mathcal{I}_i(\langle v, e_2 \rangle)$ be such that P π -intersects $P' = P_{t,e'}$. By Eq. (2) such $\langle t, e' \rangle$ exists. *Case 2:* $\langle v, e_2 \rangle \in \mathcal{P}_i^B$ (i.e., P_2 is of type B). Let $P' = P_{t,e'}$ be some type B path for $\langle t, e' \rangle \in \mathcal{I}_i(\langle v, e_2 \rangle) \setminus \mathcal{P}_i^A$. By Eq. (3) such a pair $\langle t, e' \rangle$ exists. By the definition of type B, P_2 does not π -intersect P' and vice-versa. Note that in either case, P' satisfies part (a) of the lemma. To prove part (b), let $w = \text{LCA}(v, t)$. Since $\langle t, e' \rangle \in \mathcal{I}^\mathcal{A}(\langle v, e_2 \rangle)$ (i.e., $e' \not\prec e_2$), it holds that w is not below x_2 . In addition, since P_1 and P_2 are new-ending $s - v$ paths ending with a distinct edge, by Cl. 4.5, it holds that d_2 , the unique divergence point of P_2 and $\pi(s, v)$, occurs on the segment $\pi(y_1, x_2)$.

Claim 4.9. *There exists an $s - v$ replacement-path protecting against e_2 , $P_3 \subseteq G \setminus \{e_2\}$, whose unique divergence point from $\pi(s, v)$ is not below w (see Fig. 5).*

Proof: First consider the case where $\langle v, e_2 \rangle \in \mathcal{P}_i^A$, see Fig. 5(a). In this case, by the selection of P' , it holds that P_2 π -intersects P' . Let $w' \in (V(\pi(w, t)) \setminus \{w\}) \cap V(P_2)$ and define $P_3 = \pi(s, w') \circ P_2[w', v]$. First, observe that w is the unique divergence point of P_3 and $\pi(s, v)$ since $P_2[w', v] \subseteq D(P_2)$. Next, observe that $e_2 \notin P_3$. This holds since $\pi(s, w') = \pi(s, w) \circ \pi(w, w')$. Since $e_2 \in \pi(w, v)$ and $E(\pi(w, w')) \cap E(\pi(w, v)) = \emptyset$, indeed the failing edge is not on P_3 . Finally, by the optimality of the BFS tree T_0 , $|P_3| = |P_2|$. Hence, the path P_3 satisfies the desired property as it diverges from $\pi(s, v)$ at w .

It remains to consider the case where $\langle v, e_2 \rangle \in \mathcal{P}_i^B$. See Fig. 5(b). Since both P_2 and P' are of type B, P_2 does not π -intersect P' and vice-versa, and hence

$$V(\pi(w, v) \cap P') \setminus \{w\} = \emptyset \text{ and also } V((\pi(w, t) \cap P_2)) \setminus \{w\} = \emptyset. \quad (7)$$

Let $w' \notin \{d(P_2), d(P'), v, t\}$ be a common point of the detours $D(P_2)$ and $D(P')$. Since $\langle t, e' \rangle \in \mathcal{I}^\mathcal{A}(P_2)$, by Eq. (1), such vertex w' exists. Let $P_3 = P'[s, w'] \circ P_2[w', v]$. We first claim that P_3 has a unique divergence point from $\pi(s, v)$ which is not below w . Let $d(P')$ be the unique divergence point of P' from $\pi(s, t)$ (which exists by Cl. 4.4(1)). Clearly, $P'[s, w'] = \pi(s, d(P')) \circ P'[d(P'), w']$. Since $P'[d(P'), w'] \subseteq D(P')$, it holds that $(P'[d(P'), w'] \cap \pi(s, w)) \setminus \{d(P')\} = \emptyset$. Since P' does not π -intersect with P_2 , by Eq. (7), it also holds that $(P'[d(P'), w'] \cap \pi(w, v)) \setminus \{d(P'), w\} = \emptyset$, and since $P_2[w', v] \subseteq D(P_2)$, overall it holds that $V(P_3[d(P'), v]) \cap V(\pi(s, v)) \setminus \{d(P'), v\} = \emptyset$.

Note that the last point common to P' and $\pi(s, v)$ is not below w and hence the unique divergence point of P_3 and $\pi(s, v)$ is not below w . In addition, observe that $P_3 \subseteq G \setminus \{e_2\}$

since $e_2 \in \pi(w, v)$ and P' does not intersect $\pi(w, v) \setminus \{w\}$. It remains to bound the length of P_3 . By Eq. (7), $P'[s, w'], P_2[s, w'] \subseteq G \setminus \{e', e_2\}$, and by the optimality of P' and P_2 , it holds that $|P'[s, w']| = |P_2[s, w']|$. The claim follows. ■

Since Algorithm **Pcons** attempt to select the replacement-path whose divergence point is as close to s as possible, (see Cl. 4.4(2)), it holds that d_2 is not below w . Altogether, w is above e_2 but not above d_2 , implying that $w \in \pi(y_1, x_2)$ as well, thus proving part (b) of the lemma. ■




Figure 5: Schematic Illustration for Lemma 4.8: (a) P_2 of type A. (b) P_2 of type B that does not intersect $\pi(w, t)$.

We conclude the analysis of Phase (S1) by showing that $\text{LastE}(P_{v,e}) \in H$ for every $\langle v, e \rangle \in \mathcal{I}_1 \setminus \mathcal{P}^C$. The high level idea of the proof is to use Lemma 4.8 to show that the existence of at least one H -new-ending path $P_{v,e}$ where $\langle v, e \rangle \in \mathcal{I}_1 \setminus \mathcal{P}^C$ implies that T_0 has expansion at least n^ϵ , so after $O(1/\epsilon)$ steps of expansion, it covers more than n vertices, leading to contradiction.

Lemma 4.10. $(\mathcal{I}_1 \setminus \mathcal{P}^C) \cap \mathcal{UP}(H) = \emptyset$.

Proof: Assume, towards contradiction, that there exists at least one uncovered pair $\langle v, e \rangle \in \mathcal{I}_1 \setminus \mathcal{P}^C$ such that $\text{LastE}(P_{v,e}) \notin H$. Let \mathcal{O} be the collection of all ordered pairs of vertices (x, y) where x is an ancestor of y in T_0 , that is, $\mathcal{O} = \{(x, y) \in V \times V \mid \text{LCA}(x, y) = x\}$. For every pair of vertices $(x, y) \in \mathcal{O}$ and index $t \in \{1, \dots, K_\epsilon\}$, define the collection of $s - y$ replacement paths in $\mathcal{P}_t^A \cup \mathcal{P}_t^B$ protecting the edges on $\pi(x, y)$ by

$$\mathcal{P}_t(x, y) = \{\langle y, e \rangle \in \mathcal{P}_t^A \cup \mathcal{P}_t^B \mid e \in \pi(x, y)\} . \quad (8)$$

The structure of our reasoning is as follows. For every $i \in \{1, \dots, K_\epsilon\}$, we define a collection $\widehat{\Pi}_i \subseteq \mathcal{O}$ of $\Omega(n^{\epsilon(i-1)})$ ordered vertex-pairs (x, y) , such that the paths $\pi(x, y)$ are of length $\Omega(n^\epsilon)$, and the internal segments of the paths $\pi(x, y)$ and $\pi(x', y')$ are vertex disjoint

for every two distinct pairs $(x, y), (x', y') \in \widehat{\Pi}_i$. Hence, overall, the total number of vertices occupied by the tree-paths connecting the pairs of $\widehat{\Pi}_i$ is $\Omega(n^{\epsilon^i})$. Solving for $i = K_\epsilon$, we get that the graph contains $\Omega(n^{1+\epsilon})$ distinct vertices and hence leading to contradiction.

We next define the sets $\widehat{\Pi}_i$ and show that $\widehat{\Pi}_i$ satisfies the following properties for every $i \in \{1, \dots, K_\epsilon\}$.

$$(Q1) \quad |\widehat{\Pi}_i| = \Omega(n^{\epsilon^{(i-1)}}).$$

(Q2) $T_0(z)$ and $T_0(z')$ are vertex-disjoint, where z, z' is the second vertex on $\pi(x, y)$ and $\pi(x', y')$ respectively, for every $(x, y), (x', y') \in \widehat{\Pi}_i$ and $T_0(z), T_0(z') \subseteq T_0$ are the subtrees of T_0 rooted at z, z' respectively.

(Q3) For every $(x, y) \in \widehat{\Pi}_i$, $\mathcal{P}_{K_\epsilon-i+1}(x, y)$ contains at least n^ϵ pairs whose corresponding $s - y$ replacement paths end with a distinct last edge.

$$(Q4) \quad |\pi(x, y)| \geq n^\epsilon \text{ for every } (x, y) \in \widehat{\Pi}_i.$$

We now construct $\widehat{\Pi}_i$ inductively and show by induction that it satisfies these properties. For $i = 1$, let $\widehat{\Pi}_1 = \{(s, v)\}$ where v is the vertex satisfying that there exists an $e \in \pi(s, v)$ such that $\langle v, e \rangle \in \mathcal{I}_1 \setminus \mathcal{P}^C$ and $P_{v,e}$ is H -new-ending (i.e., $\langle v, e \rangle \in \mathcal{UP}(H) \cap (\mathcal{I}_1 \setminus \mathcal{P}^C)$).

Properties (Q1-Q2) hold vacuously as $\widehat{\Pi}_1$ contains (only) one pair. We now verify (Q3), that is, we show that $\mathcal{P}_{K_\epsilon}(x, y)$ contains pairs corresponding to at least n^ϵ replacement paths whose last edges are distinct.

Since $\langle v, e \rangle \in \mathcal{I}_1 \setminus \mathcal{P}^C$, it follows that $\langle v, e \rangle \notin \mathcal{P}_{K_\epsilon}^C$ and hence $P_{v,e}$ is of type A or B with respect to the collection \mathcal{P}_{K_ϵ} . Let $J \in \{A, B\}$, be such that $\langle v, e \rangle \in \mathcal{P}_{K_\epsilon}^J$ (so, $P_{v,e}$ is of type (J)). Recall that $\vec{\mathcal{P}}_{K_\epsilon}^J(v)$ is the collection of all v 's pairs in $\mathcal{P}_{K_\epsilon}^J$ ordered in increasing distance of their failing edge from v . Recall that the algorithm adds $\lceil n^\epsilon \rceil$ distinct last edges of the paths corresponding to the *first* pairs in this ordering into H . Hence, by the fact that $\text{LastE}(P_{v,e})$ was not added to H , it follows that the paths of the pairs of $\mathcal{P}_{K_\epsilon}^J(v)$ end with more than $\lceil n^\epsilon \rceil$ distinct edges, hence (Q3) holds.

Finally, to see (Q4), observe that the replacement path of every pair $\langle v, e \rangle \in \mathcal{P}_{K_\epsilon}(s, v)$ protects a different edge on $\pi(s, v)$, i.e., $e \neq e'$ for every $\langle v, e \rangle, \langle v, e' \rangle \in \mathcal{P}_{K_\epsilon}(s, v)$. In addition, $e_j \in \pi(s, v)$ for every pair $\langle v, e_j \rangle \in \mathcal{P}_{K_\epsilon}(s, v)$. Hence, $|\pi(s, v)| \geq |\mathcal{P}_{K_\epsilon}(s, v)| > n^\epsilon$, as required, and (Q4) holds.

For the inductive step, assume that the collection $\widehat{\Pi}_{i-1}$ is given and satisfies (Q1-Q4). We now describe the construction of $\widehat{\Pi}_i$ and show that it satisfies the properties as well. To do that, every pair $(x, y) \in \widehat{\Pi}_{i-1}$ is used to produce $\Omega(n^\epsilon)$ new pairs $(x_1, y_1), \dots, (x_k, y_k)$ for $k \geq \lceil n^\epsilon \rceil$ in the following manner. Let $\mathcal{P}_{K_\epsilon-i+2}^{UN}(x, y)$ be a maximum collection of replacement *paths* corresponding to the pairs of $\mathcal{P}_{K_\epsilon-i+2}(x, y)$ each ending with a distinct last edge. By the induction assumption for Property (Q3), $|\mathcal{P}_{K_\epsilon-i+2}^{UN}(x, y)| \geq n^\epsilon$. Let $E_{i-1}(x, y) = \{e \in \pi(x, y) \mid P_{v,e} \in \mathcal{P}_{K_\epsilon-i+2}^{UN}(x, y)\}$ be the set of edges on $\pi(x, y)$ protected by the paths of $\mathcal{P}_{K_\epsilon-i+2}^{UN}(x, y)$. Let $\vec{E}_{i-1}(x, y) = \{e_{j_1}, \dots, e_{j_k}\}$ be the set of $E_{i-1}(x, y)$ ordered in increasing distance from y . As $|E_{i-1}(x, y)| = |\mathcal{P}_{K_\epsilon-i+2}^{UN}(x, y)|$, it follows that $k \geq \lceil n^\epsilon \rceil$. Let $e_{j_\ell} = (a_{j_\ell}, b_{j_\ell})$ for every $\ell \in \{1, \dots, k\}$. Since $\text{LastE}(P_1) \neq \text{LastE}(P_2)$ for every $P_1, P_2 \in \mathcal{P}_{K_\epsilon-i+2}^{UN}(x, y)$, we

can safely apply Lemma 4.8. By this lemma, for every $\ell \in \{1, \dots, k-1\}$, there exists a vertex v_ℓ and an $s-v_\ell$ replacement-path $P'_\ell = P_{v_\ell, e'_\ell}$ for $\langle v_\ell, e'_\ell \rangle \in \mathcal{P}_{K_\epsilon-i+1}^A \cup \mathcal{P}_{K_\epsilon-i+1}^B$ one level up, such that $\text{LCA}(v_\ell, y)$ is located on $\pi(s, y)$ between the two failing edges $e_{j_\ell}, e_{j_{\ell+1}} \in \pi(x, y)$, that is, $\text{LCA}(y, v_\ell) \in \pi(a_{j_\ell}, b_{j_{\ell+1}})$. Let $\widehat{\Pi}(x, y) = \{(\text{LCA}(y, v_\ell), v_\ell) \mid 1 \leq \ell \leq k\}$ and define $\widehat{\Pi}_i = \bigcup \widehat{\Pi}(x, y)$, where the union is over all $(x, y) \in \widehat{\Pi}_{i-1}$. See Fig. 6 for an illustration for the case where $i = 2$.

We now show that $\widehat{\Pi}_i$ satisfies (Q1-Q4). By induction assumption (Q1) for $i-1$, $\widehat{\Pi}_{i-1}$ contains $\Omega(n^{\epsilon(i-2)})$ pairs, and by the construction, each pair (x, y) gives rise to a collection of pairs $\widehat{\Pi}(x, y)$ of size $\Omega(n^\epsilon)$. We now show that these pairs are distinct, and hence the claim holds. By construction, $z_1 \neq z'_1$ for every $(z_1, z_2), (z'_1, z'_2) \in \widehat{\Pi}(x, y)$, as each such vertex is located between two consecutive failing edges on $\pi(x, y)$. Note that for every $(z_1, z_2) \in \widehat{\Pi}(x, y)$, it holds that z_1 occurs on $\pi(x, y)$ strictly below the first edge on $\pi(x, y)$ (i.e., z_1 is between two failing edges on $\pi(x, y)$). Thus $z_1 \in T_0(z)$ where z is the second vertex on $\pi(x, y)$. By Property (Q2) for $\widehat{\Pi}_{i-1}$, also $z_1 \neq z'_1$ for every $(z_1, z_2) \in \widehat{\Pi}(x, y)$ and $(z'_1, z'_2) \in \widehat{\Pi}(x', y')$ for every $(x, y), (x', y') \in \widehat{\Pi}_{i-1}$. Hence, Property (Q1) holds. By a similar argument, (Q2) holds as well. In particular, by (Q2) for $\widehat{\Pi}_{i-1}$ the claim holds for $(z_1, z_2) \in \widehat{\Pi}(x, y)$ and $(z'_1, z'_2) \in \widehat{\Pi}(x', y')$ corresponding to distinct pairs $(x, y), (x', y') \in \widehat{\Pi}_{i-1}$. In addition, by the definition of $\widehat{\Pi}(x, y)$, the first vertex z_1 of each pair $(z_1, z_2) \in \widehat{\Pi}(x, y)$ is located between two different consecutive edges on $\pi(x, y)$.

We now turn to consider (Q3). By the construction of $\widehat{\Pi}(x, y)$ for $(x, y) \in \widehat{\Pi}_{i-1}$, the first vertex of every pair $(z_{j_1}, z_{j_2}) \in \widehat{\Pi}(x, y)$ is $z_{j_1} = \text{LCA}(y, z_{j_2})$ and there exists a pair of edges $e \in \pi(x, y)$ and $e' \in \pi(s, z_{j_2})$ satisfying that $e \not\prec e'$ and $\langle z_{j_2}, e' \rangle \in \mathcal{P}_{K_\epsilon-i+1}^A \cup \mathcal{P}_{K_\epsilon-i+1}^B$. Since $e \not\prec e'$, it follows that $e' \in \pi(z_{j_1}, z_{j_2})$ (i.e., e' occurs below $\text{LCA}(y, z_{j_2})$). Let $J \in \{A, B\}$ be such that $\langle z_{j_2}, e' \rangle \in \mathcal{P}_{K_\epsilon-i+2} \cap \mathcal{P}_{K_\epsilon-i+1}^J$.

Recall that the algorithm adds, at the end of step $K_\epsilon - i + 1$, $\lceil n^\epsilon \rceil$ last distinct edges of $s - z_{j_2}$ replacement paths corresponding to the first pairs in the ordered set $\vec{\mathcal{P}}_{K_\epsilon-i+1}^J(z_{j_2})$. Hence, the fact that $\langle z_{j_2}, e' \rangle \in \mathcal{P}_{K_\epsilon-i+2}^J$, implies that the last edge of $P_{z_{j_2}, e'}$ was *not* taken into H , and thus there are at least $\lceil n^\epsilon \rceil$ pairs such that each their corresponding replacement paths ending with a distinct last edge, and these pairs precede it in the ordering, i.e., their corresponding paths protect edges on $\pi(s, z_{j_2})$ *below* e' . Since all the protected edges of the pairs of $\mathcal{P}_{K_\epsilon-i+1}(z_{j_1}, z_{j_2})$ belong to the segment $\pi(z_{j_1}, z_{j_2})$, property (Q3) follows. Finally, since each replacement-path corresponds to a pair in $\mathcal{P}_{K_\epsilon-i+1}(z_{j_1}, z_{j_2})$ protecting against the failing of a distinct edge on $\pi(z_{j_1}, z_{j_2})$, Property (Q4) holds as well. The induction step holds.

We now complete the proof. By Property (Q1), $\widehat{\Pi}_{K_\epsilon}$ contains $\Omega(n^{\epsilon(K_\epsilon-1)})$ pairs. By (Q2), the paths $\pi(x, y)$ and $\pi(x', y')$ are vertex-disjoint for every $(x, y), (x', y') \in \widehat{\Pi}_{K_\epsilon}$. By (Q4), the length of each path $\pi(x, y)$ is $\Omega(n^\epsilon)$ for every $(x, y) \in \widehat{\Pi}_{K_\epsilon}$, so overall there are $\Omega(n^{\epsilon K_\epsilon}) = \Omega(n^{1+\epsilon})$ vertices in these paths, in contradiction to the fact that the number of vertices in T_0 is bounded by n . The claim follows. \blacksquare




Figure 6: Illustration for Lemma 4.10. Shown is the construction of $\widehat{\Pi}_2$ where $\widehat{\Pi}_1 = \{(s, v)\}$. By the induction base, $\pi(s, v)$ contains at least n^ϵ edges protected by the paths of type A and B whose corresponding pairs are in $\mathcal{P}_{K_\epsilon}(s, v)$. These edges are sorted in increasing distance from v . By Lemma 4.8, between any two consecutive edges (b_{j_ℓ}, b_{j_ℓ}) and $(a_{j_{\ell+1}}, b_{j_{\ell+1}})$, there is a least common ancestor $\text{LCA}(v, v_\ell)$. The length of the segment $\pi(\text{LCA}(v, v_\ell), v_\ell)$ is at least n^ϵ since the last edges of the paths protecting the n^ϵ edges closest to v_ℓ were taken into H in step 2 of Phase (S1), but the last edge of the path P_{s,v,e'_ℓ} was not taken. Hence, after $i = 2$, the subtree rooted at v contains $\Omega(n^{2^\epsilon})$ distinct vertices.

4.3.2 Analysis of Phase (S2)

We begin by showing the following.

Observation 4.11. \mathcal{P}_i^C is a (\sim) -set for every $i \in \{1, \dots, K_\epsilon\}$.

Proof: To prove this, we consider two pairs $\langle v, e \rangle, \langle t, e' \rangle$ in \mathcal{P}_i^C for some $i \in \{1, \dots, \widehat{\epsilon}\}$ and show that $\langle t, e' \rangle \notin \mathcal{I}^\sim(\langle v, e \rangle)$. By definition, $\langle v, e \rangle, \langle t, e' \rangle \in \mathcal{P}_i^C \subseteq \mathcal{P}_i$. If $\mathcal{I}^\sim(\langle v, e \rangle) \cap \mathcal{P}_i = \emptyset$ then the claim holds vacuously. So consider the remaining case. Since $P_1 = P_{v,e}$ is of type C with respect to \mathcal{P}_i , Eq. (2) and (3) imply that $\mathcal{I}^\sim(\langle v, e \rangle) \cap \mathcal{P}_i \subseteq \mathcal{P}_i^A$. Since $\langle t, e' \rangle \notin \mathcal{P}_i^A$, it holds that $\langle t, e' \rangle \notin \mathcal{I}^\sim(\langle v, e \rangle) \cap \mathcal{P}_i$, and as $\langle t, e' \rangle \in \mathcal{P}_i$, we conclude that $\langle t, e' \rangle \notin \mathcal{I}^\sim(\langle v, e \rangle)$ (by symmetry, $\langle v, e \rangle \notin \mathcal{I}^\sim(\langle t, e' \rangle)$ holds as well). The observation follows. ■

Recall that $E^-(\mathcal{TD})$ is the collection of glue edges, namely, T_0 edges that do not appear by the paths ψ of the tree-decomposition \mathcal{TD} . Sub-Phase (S2.2) and Obs. 2.2 imply:

Claim 4.12. Every glue edge $e \in E^-(\mathcal{TD})$ is protected by H .

Hence, it remains to bound the number of unprotected edges on the paths of \mathcal{TD} . The following definitions are useful in our reasoning. For a vertex v and an $s_\psi - t_\psi$ path $\psi \in \mathcal{TD}$, define

$$\mathcal{P}_{\text{miss}}(\psi, v) = \{\langle v, e \rangle \in \mathcal{P}_{\text{miss}} \mid e \in \pi(s, v) \cap \psi\},$$

as the set of uncovered pairs in H that belong to \mathcal{P} and whose corresponding replacement paths protect against the failure of the edges on $\pi(s, v) \cap \psi$. Let $E_{\text{miss}}(\mathcal{P}, \psi, v) = \{e \mid \langle v, e \rangle \in \mathcal{P}_{\text{miss}}(\psi, v)\}$ be the corresponding last-unprotected edges by H on $\pi(s, v) \cap \psi$ and let $E_{\text{miss}}(\mathcal{P}, \psi) = \bigcup_{v \in V} E_{\text{miss}}(\mathcal{P}, \psi, v)$. By Cl. 4.12,

$$E_{\text{miss}}(\mathcal{P}) = \bigcup_{\psi \in \mathcal{TD}} E_{\text{miss}}(\mathcal{P}, \psi).$$

Note that the replacement paths of the pairs of $\mathcal{P}_{\text{miss}}(\psi, v)$ may end with the same last edge. We now identify a set $\mathcal{P}_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v)$ of unique representatives for each last edge as follows. For every edge e' that has several replacement paths $P_{v,e}$ for $\langle v, e \rangle \in \mathcal{P}_{\text{miss}}(\psi, v)$ whose last edge $\text{LastE}(P_{v,e}) = e'$, we pick one representative pair $\langle v, e^* \rangle$ corresponding to the path P_{v,e^*} whose failing edge e^* is closest to s among all other candidates. Formally, let $\mathcal{LE}_{\text{miss}}(\mathcal{P}, \psi, v) = \{\text{LastE}(P_{v,e}) \mid \langle v, e \rangle \in \mathcal{P}_{\text{miss}}(\psi, v)\}$ be the last edges of the replacement paths of the pairs in $\mathcal{P}_{\text{miss}}(\psi, v)$. For every $e' \in \mathcal{LE}_{\text{miss}}(\mathcal{P}, \psi, v)$, let $\mathcal{P}(e', \mathcal{P}, \psi, v) = \{\langle v, e \rangle \in \mathcal{P}_{\text{miss}}(\psi, v) \mid \text{LastE}(P_{v,e}) = e'\}$. The representative pair for the last edge e' denoted by $\hat{P}(e') = \langle v, e^* \rangle$ for $\langle v, e^* \rangle \in \mathcal{P}(e', \mathcal{P}, \psi, v)$ satisfying that $\text{dist}(s, e^*, \pi(s, v)) < \text{dist}(s, e'', \pi(s, v))$ for every $e'' \neq e^*$ and $\langle v, e'' \rangle \in \mathcal{P}(e', \mathcal{P}, \psi, v)$. Finally, define $\mathcal{P}_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v) = \{\hat{P}(e') \mid e' \in \mathcal{LE}_{\text{miss}}(\mathcal{P}, \psi, v)\}$ and

$$E_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v) = \{e_j \mid \langle v, e_j \rangle \in \mathcal{P}_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v)\}.$$

We proceed by showing that $E_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v)$ is either *empty* or sufficiently *large*.

Lemma 4.13. *For every $v \in V$ and every $\psi \in \mathcal{TD}$, if $E_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v) \neq \emptyset$, then:*

(a) $|E_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v)| \geq \lceil n^\epsilon \rceil$, and (b) *the first $\lceil n^\epsilon \rceil$ edges in $E_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v)$ are contained in $\pi_{j^*}(s, v) \subseteq \pi(s, v)$, which is the highest heavy subsegment that intersects ψ with respect to \mathcal{P} .*

Proof: Recall that in Sub-Phase (S2.1), the $s - v$ shortest-path $\pi(s, v)$ was partitioned into $k' = \lfloor \log |\pi(s, v)| \rfloor$ segments where the j 'th segment $\pi_j(s, v)$ is given by $\pi_j(s, v) = \pi(u_{i_{j-1}}, u_{i_j})$ for every $j \in \{1, \dots, k'\}$. Let $e_1 \in \pi(s, v)$ be the closest edge to s in $E_{\text{miss}}(\mathcal{P}, \psi, v)$ and $e_1 \in E_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v)$. Since $e_1 \in E_{\text{miss}}(\mathcal{P}, \psi, v)$, it implies that e_1 belongs to a heavy subsegment and in particular, $e_1 \in \pi_{j^*}(s, v)$. First, consider the case where this subsegment is fully contained in ψ , i.e., that $\pi_{j^*}(s, v) \subseteq \psi$. Recall that $\mathcal{P}_{j^*}(v) = \{\langle v, e \rangle \in \mathcal{P} \mid e \in \pi_{j^*}(s, v)\}$ is the collection of $s - v$ replacement paths in \mathcal{P} that protect the edges on $\pi_{j^*}(s, v)$ and $\mathcal{LE}_{j^*}(v)$ is the corresponding last edges of these paths. Since $e_1 \in E_{\text{miss}}(\mathcal{P}, \psi, v)$, it follows that $\pi_{j^*}(s, v)$ is heavy with respect to \mathcal{P} , i.e., that $|\mathcal{LE}_{j^*}(v)| > \lceil n^\epsilon \rceil$. Hence, the pairs of $\mathcal{P}_{j^*}(v)$ correspond to at least $\lceil n^\epsilon \rceil$ replacement paths that end with distinct last edges, implying that $|E_{\text{miss}}^{\text{UN}}(\mathcal{P}, \psi, v)| \geq \lceil n^\epsilon \rceil$.

Next, consider the complementary case where $\pi_{j^*}(s, v) \not\subseteq \psi$. This implies that e_1 belongs to the subsegment $\pi_U(\psi, v)$ or $\pi_L(\psi, v)$ that intersects with ψ . Assume first that $e_1 \in \pi_U(\psi, v)$. Since $e_1 \in E_{\text{miss}}(\mathcal{P}, \psi, v)$, it follows that $|\mathcal{LE}_U(\mathcal{P}, \psi, v)| \geq \lceil n^\epsilon \rceil$ and hence $\mathcal{P}_U(\psi, v)$ contains pairs corresponding to at least $\lceil n^\epsilon \rceil$ replacement paths that end with a distinct last edge. The case where $e_1 \in \pi_L(\psi, v)$ is analogous. ■

Note that if $E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v) = \emptyset$, then also $E_{miss}(\mathcal{P}, \psi, v) = \emptyset$, so this set needs not concern us anymore. Hence hereafter we concentrate on vertices v with a large set $E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v)$. For such a vertex v , let

$$\vec{E}_{miss}^{\text{UN}}(\mathcal{P}, \psi, v) = \{e_{i_1}, \dots, e_{i_\ell}\}$$

be the edges of $E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v)$ ordered in *increasing* distance from s . By Lemma 4.13, $\ell \geq \lceil n^\epsilon \rceil$. Define $\mathcal{D}(\mathcal{P}, \psi, v) = \{D(P_{v, e_{i_j}}) \mid j \in \{1, \dots, \lceil n^\epsilon \rceil\}\}$ as the collection of the detours protecting against the failure of the first $\lceil n^\epsilon \rceil$ ordered edges in the ordering $\vec{E}_{miss}^{\text{UN}}(\mathcal{P}, \psi, v)$. Note that by the definition of $E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v)$, each of the detours in $\mathcal{D}(\mathcal{P}, \psi, v)$ ends with a distinct last edge (in particular, by Cl. 4.6, these detours are vertex disjoint, except for the terminal v).

In addition, for a vertex v with a large set $E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v)$, Let $e^*(\mathcal{P}, \psi, v) \in E_{miss}(\mathcal{P}, \psi, v)$ be the closest edge to s on $\pi(s, v) \cap \psi$ among all edges in $E_{miss}(\mathcal{P}, \psi, v)$. Hence, $e^*(\mathcal{P}, \psi, v) \in E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v)$. Note that by the end of the Sub-Phases (S2.2.1-2) and by Cl. 4.5, the divergence point of $P_{v, e^*(\mathcal{P}, \psi, v)}$ from $\pi(s, v)$ must occur on ψ , i.e., $d(P_{v, e^*(\mathcal{P}, \psi, v)}) \in \pi(s, v) \cap \psi$. (This is because the last edges of the new ending paths protecting the first failing edges on each subsegment $\pi_j(s, v)$ and the intersected segments $\pi_U(s, v), \pi_L(s, v)$ were added into H , the divergence point of the H -new-ending paths protecting the other edges on these segments are internal to their segments.) Define the segments

$$\sigma(\mathcal{P}, \psi, v) = \pi(d(P_{v, e^*(\mathcal{P}, \psi, v)}), \text{LCA}(v, t_\psi)),$$

and the segment collection

$$\mathcal{SG}(\mathcal{P}, v) = \{\sigma(\mathcal{P}, \psi, v) \mid E_{miss}(\mathcal{P}, \psi, v) \neq \emptyset\}, \quad (9)$$

where ψ is an $s_\psi - t_\psi$ path, for illustration see Fig. 7.

We next claim that each of the detours of $\mathcal{D}(\mathcal{P}, \psi, v)$ is sufficiently long.

Lemma 4.14. *If $\mathcal{D}(\mathcal{P}, \psi, v)$ is nonempty, then $|D_i| \geq |\sigma(\mathcal{P}, \psi, v)|/4$ for every $D_i \in \mathcal{D}(\mathcal{P}, \psi, v)$.*

Proof: By Cl. 4.13(b), the first $\lceil n^\epsilon \rceil$ edges in $E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v)$ are contained in $\pi_{j^*}(s, v) \subseteq \pi(s, v)$, which is the highest *heavy* subsegment with respect to \mathcal{P} and v that intersects $\pi(s, v)$. Hence letting $d^* = d(P_{v, e^*(\mathcal{P}, \psi, v)})$, we get that d^* occurs on $\pi_{j^*}(s, v)$ and the detours of a nonempty set $\mathcal{D}(\mathcal{P}, \psi, v)$ protect the failing of edges on $\pi_{j^*}(s, v)$.

By the right inequality of Eq. (5), for every $D_i \in \mathcal{D}(\mathcal{P}, \psi, v)$,

$$\begin{aligned} |V(D_i)| &\geq \sum_{j' > j^*} |\pi_{j'}(s, v)| \geq |\pi_{j^*}(s, v)|/2 \geq \sum_{j' \geq j^*} |\pi_{j'}(s, v)|/4 \\ &\geq |\pi(d^*, v)|/4 = |\sigma(\mathcal{P}, \psi, v)|/4. \end{aligned} \quad (10)$$

The lemma follows. ■

Observation 4.15. $E_{miss}(\mathcal{P}, \psi) \subseteq \bigcup_{\sigma \in \mathcal{SG}(\mathcal{P}, \psi)} \sigma$.




Figure 7: Illustration of the segments $\sigma(\mathcal{P}, \psi, v)$. For brevity, we omit \mathcal{P} and simply write $\sigma(\psi, v)$. Shown is an $s_\psi - t_\psi$ path $\psi \in \mathcal{TD}$ and (a) a segment $\sigma(\psi, v)$ defined by the divergence point of the replacement-path $P_{v,e^*(\mathcal{P},\psi,v)}$ protecting the highest unprotected edge on $\pi(s,v) \cap \psi$ and the LCA of v and t_ψ ; (b) the independent segments $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$, where $\ell = \max\{|\sigma(\psi, v_1)|, |\sigma(\psi, v_2)|\}$ is the minimum spacing between two segments.

Proof: Assume towards contradiction that there exists an edge $e \in E_{miss}(\mathcal{P}, \psi) \setminus \bigcup_{\sigma \in \mathcal{SG}(\mathcal{P}, \psi)} \sigma$. Let $v \in V$ be such that $\text{LastE}(P_{v,e}) \notin H$, i.e., $e \in E_{miss}(\mathcal{P}, \psi, v)$.

Recall that $e^*(\mathcal{P}, \psi, v)$ is the closest edge to s in $E_{miss}(\mathcal{P}, \psi, v)$, hence $e^*(\mathcal{P}, \psi, v) \in E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v)$, and e is not above $e^*(\mathcal{P}, \psi, v)$ on $\pi(s, v)$. In addition, since $e \in \pi(s, v) \cap \psi$, it holds that e is above $\text{LCA}(v, t_\psi)$. Altogether, we get that $e \in \sigma(\mathcal{P}, \psi, v)$, contradiction. The observation follows. \blacksquare

From now on, we focus on a particular path ψ in the tree decomposition \mathcal{TD} . We proceed by defining a notion of independence between two segments $\sigma_i = \sigma(\mathcal{P}, \psi, v_i)$ and $\sigma_j = \sigma(\mathcal{P}, \psi, v_j)$ in $\mathcal{SG}(\mathcal{P}, \psi)$ (see Eq. (9) for the definition of $\mathcal{SG}(\mathcal{P}, \psi)$). Let x_i (resp. x_j) be the first vertex of σ_i (resp., σ_j) and let $y_i = \text{LCA}(v_i, t_\psi)$ (resp., $y_j = \text{LCA}(v_j, t_\psi)$) be the last vertex of σ_i, σ_j .

Definition 4.16 (Independent Segments). *Let $\sigma_i = \pi(x_i, y_i), \sigma_j = \pi(x_j, y_j) \in \mathcal{SG}(\mathcal{P}, \psi)$ be such that $\text{dist}(s, x_i, G) \leq \text{dist}(s, x_j, G)$ and let $\ell = \max\{|\sigma_i|, |\sigma_j|\}$. Then, σ_i and σ_j are independent if $\text{dist}(s, x_j, G) - \text{dist}(s, y_i, G) \geq \ell$, otherwise they are dependent.*

By Lemma 4.13, we have the following.

Observation 4.17. *For a vertex v with $E_{miss}^{\text{UN}}(\mathcal{P}, \psi, v) \neq \emptyset$, we have $|\sigma(\mathcal{P}, \psi, v)| = \Omega(n^\epsilon)$.*

Set $\mathcal{SG}'(\mathcal{P}, \psi) \leftarrow \mathcal{SG}(\mathcal{P}, \psi)$. We now compute a collection of maximal weighted independent set $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$ greedily by adding to $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$ at each step the segment $\sigma(\mathcal{P}, \psi, v) \in \mathcal{SG}'(\mathcal{P}, \psi)$ whose length is maximal among all remaining segments $\mathcal{SG}'(\mathcal{P}, \psi)$.

and removing from it the segments $\sigma(\mathcal{P}, \psi, v'')$ that are *dependent* with $\sigma(\mathcal{P}, \psi, v)$. For illustration see Fig. 7(b). The next observation shows that the total length of the independent set $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$ is of the same order as the original set $\mathcal{SG}(\mathcal{P}, \psi)$.

Claim 4.18. $\sum_{\sigma \in \mathcal{SG}_{IS}(\mathcal{P}, \psi)} |\sigma| \geq |E_{miss}(\mathcal{P}, \psi)|/5$ (see Fig. 8).

Proof: Note that by the definition of independence, $\sigma \cap \sigma' = \emptyset$ for every $\sigma, \sigma' \in \mathcal{SG}_{IS}(\mathcal{P}, \psi)$. By Obs. 4.15, $|\bigcup_{\sigma \in \mathcal{SG}(\mathcal{P}, \psi)} \sigma| \geq |E_{miss}(\mathcal{P}, \psi)|$. Let $\Delta(\sigma)$ be the collection of segments discarded from $\mathcal{SG}'(\mathcal{P}, \psi)$ before adding σ into $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$. We now show that

$$\left| \bigcup_{\sigma' \in \Delta(\sigma)} \sigma' \right| \leq 5 \cdot |\sigma|.$$

Note that every σ' in $\Delta(\sigma)$ is dependent with respect to σ . The inequality follows by the maximality of σ in $\Delta(\sigma)$ and by the definition of independence. The claim follows by summing over all independent segments σ in $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$. ■




Figure 8: Illustration for Claim 4.18. Shown is an $s_\psi - t_\psi$ path $\psi \in \mathcal{TD}$ and the segment σ of length $|\sigma| = \ell$ that was taken into $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$. The dependent segments $\Delta(\sigma)$ that were discarded from $\mathcal{SG}'(\mathcal{P}, \psi)$ when adding σ occupied at most $5|\sigma|$ vertices from ψ . To see this, observe that in the most extreme scenario, σ' and σ'' shown in the figure are in $\Delta(\sigma)$. By the maximality of the length of σ at the time it was taken into $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$, it holds that $|\sigma'|, |\sigma''| \leq \ell$.

We next show the following (for the given ψ and \mathcal{P}).

Lemma 4.19. $n^\epsilon \cdot \sum_{\sigma \in \mathcal{SG}_{IS}(\mathcal{P}, \psi)} |\sigma| \leq |\bigcup_{P \in \mathcal{P}_{miss}(\mathcal{P}, \psi)} V(D(P))|.$

To prove the lemma, we consider an iterative process on the set $VS = \{v \mid \sigma(\mathcal{P}, \psi, v) \in \mathcal{SG}_{IS}(\mathcal{P}, \psi)\}$, the set of vertices whose segment is in the independent set $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$. In this process, the detours of these vertices v' are added in decreasing distance of $\text{LCA}(v', t_\psi)$ and s . (Note that the order is strictly decreasing since the segments are independent and hence also vertex disjoint.) Formally, let $\overrightarrow{VS} = \{v_{i_1}, \dots, v_{i_k}\}$ be the collection of VS vertices sorted in decreasing distance of $\text{LCA}(v_{i_j}, t_\psi)$ and s , i.e., $\text{dist}(s, \text{LCA}(v_{i_1}, t_\psi)) > \dots > \text{dist}(s, \text{LCA}(v_{i_k}, t_\psi))$. Starting with $G'_1 = \emptyset$, at step $\tau \geq 1$, let

$$G'_{\tau+1} = G'_\tau \cup \bigcup_{D' \in \mathcal{D}(\mathcal{P}, \psi, v_{i_\tau})} D'.$$

Let $\widehat{G} = G'_k$ be the final subgraph. Hence,

$$\widehat{G} \subseteq \bigcup_{\langle v, e \rangle \in \mathcal{P}} D(P_{v, e}). \quad (11)$$

Lemma 4.20. *For every $\tau \in \{1, \dots, k\}$, $|V(G'_\tau) \setminus V(G'_{\tau-1})| \geq n^\epsilon/4 \cdot |\sigma(\mathcal{P}, \psi, v_{i_\tau})|$.*

Proof: Consider a specific detour $\widetilde{D} \in \mathcal{D}(\mathcal{P}, \psi, v_{i_\tau})$ of a path $P_{v_{i_\tau}, e}$ for $e \in E_{\text{miss}}(\mathcal{P}, \psi, v_{i_\tau})$. First note that since each of the detours in $\mathcal{D}(\mathcal{P}, \psi, v_{i_\tau})$ ends with a distinct edge, by Cl. 4.6(2), these detours are vertex disjoint (besides the common endpoint v_{i_τ}). We distinguish between two cases. Case (1): the internal segment of the detour \widetilde{D} does not intersect with the vertices of $G'_{\tau-1}$. In this case, by Lemma 4.14, $|V(\widetilde{D})| \geq |\sigma(\mathcal{P}, \psi, v_{i_\tau})|/4$.

Case (2): the internal segment of \widetilde{D} intersects at least one vertex of $G'_{\tau-1}$. Let w be the first internal vertex of \widetilde{D} that occurs in some D' , i.e., $w \in V(\widetilde{D}) \setminus \{d, v_{i_\tau}\}$ where d (resp., v_{i_τ}) is the first (resp., last) vertex of \widetilde{D} . That is, $\widetilde{D}[d, w] \setminus \{w, d\}$ is vertex disjoint with $G'_{\tau-1}$. Note that by the definition of the segments, $d \in \sigma(\mathcal{P}, \psi, v_{i_\tau})$. We now show that $|\widetilde{D}(d, w)| \geq |\sigma(\mathcal{P}, \psi, v_{i_\tau})|$. By Cl. 4.6(2) and by the ordering \overrightarrow{VS} , D' is a detour of some $P_{v_{i_j}, e'}$ path for v_{i_j} such that $j < \tau$. Let $\sigma_\tau = \sigma(\mathcal{P}, \psi, v_{i_\tau}) = \pi(x_\tau, y_\tau)$ and $\sigma_j = \sigma(\mathcal{P}, \psi, v_{i_j}) = \pi(x_j, y_j)$ (this definition is consistent, as the segments σ are subpaths on the tree T_0).

By the ordering of the insertion into \widehat{G} , it holds that $\text{dist}(s, y_\tau, G) < \text{dist}(s, x_j, G)$. In addition, since σ_j and σ_τ are independent (i.e., in $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$),

$$|\pi(y_\tau, x_j)| \geq \max\{|\sigma_\tau|, |\sigma_j|\}. \quad (12)$$

This case is further divided into two cases depending on whether or not the failing edge e' (protected by the detour D') occurs on $\widetilde{D}[d, w]$. First, consider the case where $e' = (x', y') \in \widetilde{D}[d, w]$, that is, e' occurs on \widetilde{D} before the first common vertex. For illustration, see Fig. 9(a). By definition, $e' \in \sigma_j$ and by the independence of the subsegments, it holds that $e' \notin \sigma_\tau$. By the ordering, y_τ occurs on ψ not below d' where d' is the first vertex of the detour D' . We get that

$$|\widetilde{D}[d, w]| \geq |\widetilde{D}[d, y']| \geq |\pi(d, y')| \geq |\pi(y_\tau, x_j)| \geq |\sigma_\tau|,$$

where the last inequality holds by Eq. (12). Finally, we turn to consider the complementary case where $e' \notin \tilde{D}[d, w]$. Then there are two $d - w$ shortest paths in $G \setminus \{e'\}$ given by $P_1 = \pi(d, d') \circ D'[d', w]$ and $P_2 = \tilde{D}[d, w]$ where $P_1 \subseteq P_{v_{ij}, e'}$ and $P_2 \subseteq P_{v_{i\tau}, e'}$. As P_1 is optimal in $G \setminus \{e'\}$,

$$|P_2| \geq |P_1| \geq |\pi(d, d')| \geq |\pi(y_\tau, x_j)| \geq |\sigma_\tau| ,$$

where the last inequality follows again by Eq. (12). ■




Figure 9: Illustration for Lemma 4.20. Shown is an $s_\psi - t_\psi$ path $\psi \in \mathcal{TD}$ and two independent segments σ_τ, σ_j in $\mathcal{SG}_{IS}(\mathcal{P}, \psi)$. The detour \tilde{D} was added to \hat{G} at step τ and it intersects D' , which was added to \hat{G} at step $j < \tau$. (a) The detour \tilde{D} traverses e' , the failing edge protected by D' , before the first mutual vertex w . (b) $e' \notin \tilde{D}[d, w]$ (where d is the first vertex of D).

By the last two lemmas we get:

Lemma 4.21. $|\bigcup_{P \in \mathcal{P}_{miss}(\mathcal{P}, \psi)} V(D(P))| = \Omega(n^\epsilon \cdot |E_{miss}(\mathcal{P}, \psi)|)$.

Corollary 4.22. $|E_{miss}(\mathcal{P})| = O(n^{1-\epsilon} \cdot \log n)$ for every $\mathcal{P} \in \{\mathcal{I}_2, \dots, \mathcal{P}_1^C, \dots, \mathcal{P}_{K_\epsilon}^C\}$.

Proof: Recall that by Cl. 4.12, the edges that are not covered by the paths of the tree-decomposition are protected by H . Hence, the set of unprotected edges is given by

$$E_{miss}(\mathcal{P}) = \bigcup_{\psi \in \mathcal{TD}} E_{miss}(\mathcal{P}, \psi) . \quad (13)$$

The recursive tree-decomposition algorithm of [2] (Procedure Partition therein) consists of $O(\log n)$ levels. Let $\psi_{i_1}, \dots, \psi_{i_\ell} \in \mathcal{TD}$ be a collection of paths constructed in the same

recursion level. Let $v_{i_1}, \dots, v_{i_\ell}$ be the first vertices of these paths respectively (the path endpoint that is closer to s). It can be shown that the subtrees $T(v_{i_1}), \dots, T(v_{i_\ell}) \subseteq T_0$ are vertex disjoint, and hence $e \not\sim e'$ for every $e \in \psi_{i_j}$ and $e' \in \psi_{i_{j'}}$. By Lemma 4.21, the number of vertices occupied by the detour segments protecting the edges of $E_{\text{miss}}(\mathcal{P}, \psi_{i_j})$ is $\Omega(n^\epsilon \cdot |E_{\text{miss}}(\mathcal{P}, \psi_{i_j})|)$. Since \mathcal{P} is a (\sim) -set, the internal detour segments protecting the edges on ψ_{i_j} and on $\psi_{i_{j'}}$ are vertex-disjoint. Note that the definition of independence between detours, refers to empty intersection of their internal segments (excluding the first and the last vertices). Since the length of each detour is $\Omega(n^\epsilon)$ (see Obs. 4.17 and the proof of Lemma 4.20), this is negligible. Overall, we get that the number of vertices occupied by these detours is bounded by $\Omega(n^\epsilon \cdot \sum_{j=1}^k |E_{\text{miss}}(\mathcal{P}, \psi_{i_j})|)$. As the number of vertices is bounded by n , we get that $\sum_{j=1}^k |E_{\text{miss}}(\mathcal{P}, \psi_{i_j})| = O(n^{1-\epsilon})$. Summing over all $O(\log n)$ recursion levels, and combining with Eq. (13),

$$|E_{\text{miss}}(\mathcal{P})| = \sum_{\psi \in \mathcal{TD}} |E_{\text{miss}}(\mathcal{P}, \psi)| = O(n^{1-\epsilon} \cdot \log n).$$

The lemma follows. \blacksquare

We are now ready to complete the proof of Thm. 3.1.

Proof: [Thm. 3.1] Recall that $\mathcal{UP} = \mathcal{I}_1 \cup \bigcup_{i=0}^{K_\epsilon} \mathcal{P}_i^C$. By Lemma 4.10, the uncovered pairs in H (i.e., pairs that correspond to H -new-ending paths) are in $\bigcup_{i=0}^{K_\epsilon} \mathcal{P}_i^C$. Finally, using Cor. 4.22 and summing over all $O(1/\epsilon)$ (\sim) -sets \mathcal{P} yields Lemma 4.7, which bounds the number of unprotected edges (that need to be reinforced) as in the theorem. Lemma 4.2 bounds the size of H , hence the number $b(n)$ of backup edges. Consequently the theorem follows. \blacksquare

5 Lower Bound

In this section, we establish lower bounds on the size of the ϵ FT-BFS structures. These bounds match the upper bound of Sec. 3 up to logarithmic factors in both the number of reinforced edges and the size of the construct. These lower bound constructions are generalizations of [14]. We first consider the single source case. Note that for $\epsilon \in [1/2, 1]$, by the lower bound FT-BFS in [14] $\Omega(n^{3/2})$ edges are required. Hence, it remains to establish the lower bound for $\epsilon \in (0, 1/2)$.

Theorem 5.1. *For every $\epsilon \in (0, 1/2)$, there exists an n -vertex graph $G(V, E)$ and a source node $s \in V$ such that any ϵ FT-BFS tree rooted at s with at most $\lfloor n^{1-\epsilon}/6 \rfloor$ reinforced edges has $\Omega(n^{1+\epsilon})$ edges. In other words, there exists a graph for which any $(b(n), r(n))$ FT-BFS structure for $r(n) = \Omega(n^{1-\epsilon})$ requires $\Omega(\min\{n^{1+\epsilon}, n^{3/2}\})$ backup edges.*

Proof: Let us first describe the structure of $G = (V, E)$. Set $d_\epsilon = \lfloor n^\epsilon/4 \rfloor$ and $k_\epsilon = \lfloor n^{1-2\epsilon} \rfloor$. We first describe the structure of a subgraph G_ϵ which provides the basic building block of the construction. In particular, the final graph G consists of k_ϵ copies of the graph G_ϵ denoted by $G_{\epsilon,1}, \dots, G_{\epsilon,k_\epsilon}$ that are connected to the source vertex s as will be described later.

We begin by describing the structure of the i th copy $G_{\epsilon,i}$ (the copies are identical), which consists of four main components. The first is a path

$$\pi_i = [s_i = v_1^i, \dots, v_{d_\epsilon+1}^i = v_i^*]$$

of length d_ϵ . The second component consists of a node set $Z_i = \{z_1^i, \dots, z_{d_\epsilon}^i\}$ and a collection of d_ϵ disjoint paths of decreasing length, $P_1^i, \dots, P_{d_\epsilon}^i$, where $P_j^i = [v_j^i = p_{j,1}^i, \dots, p_{j,t_j}^i = z_j^i]$ connects v_j^i on π_i with z_j^i and its length is $t_j = |P_j^i| = 6 + 2(d_\epsilon - j)$, for every $j \in \{1, \dots, d_\epsilon\}$. Altogether, the set of nodes in these paths,

$Q_i = \bigcup_{j=1}^{d_\epsilon} V(P_j^i)$, is of size $|Q_i| = d_\epsilon^2 + 5d_\epsilon = \Theta(n^{2\epsilon})$. The third component is a set of nodes X_i of size $n - 1 - k_\epsilon(|\pi_i| + |Q_i| + |Z_i|) = \Theta(n^{2\epsilon})$, all connected to the terminal node v_i^* . The last component is a *complete* bipartite graph $B_i = (X_i, Z_i)$ connecting every vertex in X_i to every vertex in Z_i .

So far, $V(G_{\epsilon,i}) = V(\pi_i) \cup Z_i \cup X_i \cup Q_i$ and

$$E(G_{\epsilon,i}) = E(\pi_i) \cup \{(v_i^*, x_t^i) \mid x_t^i \in X_i\} \cup E(B_i) \cup \bigcup_{j=1}^{d_\epsilon} E(P_j^i).$$

Finally, s is connected via a star to the first vertex s_i of each path π_i for every $i \in \{1, \dots, k_\epsilon\}$, see Fig. 10(b) for illustration.

Overall, $V(G) = \{s\} \cup \bigcup_{i=1}^{k_\epsilon} V(G_{\epsilon,i})$ and $E(G) = E(s) \cup \bigcup_{i=1}^{k_\epsilon} V(G_{\epsilon,i})$, where $E(s) = \{(s, s_i) \mid i \in \{1, \dots, k_\epsilon\}\}$.

A BFS tree T_0 rooted at s for this G is given by

$$E(T_0) = E(s) \cup \bigcup_{i=1}^{k_\epsilon} \left(\pi_i \cup \bigcup_{x \in X_i} (v_i^*, x) \cup \bigcup_{t=1}^{k_\epsilon} (x_t^i, z_t^i) \cup \bigcup_{j=1}^{d_\epsilon} E(P_j^i) \setminus \{(p_{j,\ell_j}^i, p_{j,\ell_j-1}^i)\} \right),$$

where $\ell_j = t_j - (d_\epsilon - j)$ for every $j \in \{1, \dots, d_\epsilon\}$.

Let $B = \bigcup_{i=1}^{k_\epsilon} B_i$ be the collection of edges on the k_ϵ complete bipartite graphs B_i and let $\Pi = \bigcup_{i=1}^{k_\epsilon} \pi_i$.

Observation 5.2. (a) $|V(G_{\epsilon,i})| = n/k_\epsilon - 1$ and hence $|V(G)| = n$.

(b) $|E(G)| = \Omega(|EB|) = \Omega(n^{1+\epsilon})$.

(c) $|E(\Pi)| = k_\epsilon \cdot d_\epsilon \geq \lfloor n^{1-\epsilon}/5 \rfloor$.

We now show that every ϵ FT-BFS structure H must contain a constant fraction of the edges in $B = \bigcup_{i=1}^{k_\epsilon} B_i$, namely, the edges $e_{t,j}^i = (x_t^i, z_j^i)$ (the thick edges in Fig. 10(b)).

Note that the edges of Π are “costly” in the sense that to protecting against their failure, many edges connecting X_i and Z_i should be introduced into the fault-tolerant structure

H . To provide a succinct structure, one can choose to fortify these edges, however, since $|E(\Pi)| = \Omega(n^{1-\epsilon})$, upon setting the fortification budget as in the statement of Thm. 5.1, a constant fraction of these edges could not be reinforced, resulting eventually in a dense structure. We now formalize this intuition.

For ease of analysis, consider a partition the edges of B into $d_\epsilon \cdot k_\epsilon$ disjoint subsets corresponding to the number of edges in Π . Let $e_j^i = (v_j^i, v_{j+1}^i) \in \pi_i$ be the j 'th edge on the path $\pi_i \subseteq G_{\epsilon,i}$ and define $E_j^i = \{(x_t^i, z_j^i) \mid x_t^i \in X_i\}$. Observe that E_j^i and $E_{j'}^{i'}$ are disjoint for every pair $(i, j) \neq (i', j')$ for every $i \in \{1, \dots, k_\epsilon\}$ and every $j \in \{1, \dots, d_\epsilon\}$ and in addition, $B = \bigcup_{i=1}^{k_\epsilon} \bigcup_{j=1}^{d_\epsilon} E_j^i$.

For every ϵ FT-BFS structure $H \subseteq G$, let $E'(H)$ be the set of at most $\lfloor n^{1-\epsilon}/6 \rfloor = O(n^\epsilon)$ reinforced edges in H (which is the maximum allowed number of reinforced edges by the statement of the theorem).

Claim 5.3. $E_j^i \subseteq H$ for every $e_j^i \in \Pi \setminus E'(H)$.

Proof: Assume, towards contradiction, that H does not contain $e' = (x_t^i, z_j^i) \in E_j^i$ for some $x_t^i \in X_i$ (the bold dashed edge (x_t^i, z_j^i) in Fig. 10(b)). Note that upon failure of the edge $e_j^i = (v_j^i, v_{j+1}^i) \in \pi_i$, the unique $s - x_t^i$ shortest path connecting s and x_t^i in $G \setminus \{e_j^i\}$ is $P'_{i,j} = \pi[s, v_j^i] \circ P_j^i \circ e'$, and all other alternatives are strictly longer. Since $e' \notin H$, also $P'_{i,j} \not\subseteq H$, and therefore $\text{dist}(s, x_t^i, G \setminus \{e_j^i\}) < \text{dist}(s, x_t^i, H \setminus \{e_j^i\})$, in contradiction to the fact that H is an ϵ FT-BFS structure. It follows $E_j^i \subseteq H$. ■

We are now ready to complete of Theorem 5.1. By Obs. 5.2(c), $|E(\Pi)| \geq \lfloor n^{1-\epsilon}/5 \rfloor$. By the statement of Thm. 5.1, the ϵ FT-BFS structure contains at most $\lfloor n^\epsilon/6 \rfloor$ reinforced edges $E'(H)$. Hence, even if all those edges are taken from the edges of Π (i.e., $E'(H) \subseteq E(\Pi)$), there are still $|E(\Pi) \setminus E'(H)| = \Omega(n^{1-\epsilon})$ edges in Π that remain unenforced. By Cl. 5.3, each such edge requires that E_j^i would be added to H . Note that $|E_j^i| = |X_i| = \Omega(n^{2\epsilon})$. By the disjointness of the E_j^i sets, $|E(H)| \geq \Omega(n^{2\epsilon}) \cdot |\Pi \setminus E'(H)| = \Omega(n^{1+\epsilon})$, as required. ■

Multiple Sources. In this subsection, we consider an intermediate setting where it is necessary to construct an ϵ fault tolerant subgraph containing several ϵ FT-BFS trees in parallel, one for each source $s \in S$, for some $S \subseteq V$.

For a given subset of \mathcal{K} sources $S \subseteq V$, a subgraph $H \subseteq G$ is an ϵ FT-MBFS structure if there exists a subset $E' \subseteq E$ of $O(\mathcal{K}^\epsilon \cdot n^{1-\epsilon})$ edges such that $\text{dist}(s, v, H \setminus \{e\}) = \text{dist}(s, v, G \setminus \{e\})$ for every $s, v \in S \times V$ and every $e \in E \setminus E'$. Towards the end of this section, we show the following.

Theorem 5.4. *For every real $\epsilon \in (0, 1/2]$ and $\mathcal{K} \in \{1, \dots, n\}$, there exists a graph $G = (V, E)$, a subset of \mathcal{K} sources $S \subseteq V$, such that any ϵ FT-MBFS structure with at most $\lfloor (\mathcal{K}^\epsilon \cdot n^{1-\epsilon})/6 \rfloor$ reinforced edges, contains $\Omega(\mathcal{K}^{1-\epsilon} \cdot n^{1+\epsilon})$ edges.*

Set $d_{\epsilon, \mathcal{K}} = \lfloor (n/4\mathcal{K})^\epsilon \rfloor$ and $k_{\epsilon, \mathcal{K}} = \lfloor (n/\mathcal{K})^{1-2\epsilon} \rfloor$.

We first describe the structure of a subgraph $G_{\epsilon, \mathcal{K}}$ which is a subgraph of G_ϵ in the construction for the single source case. The final constructs uses $\mathcal{K} \cdot k_{\epsilon, \mathcal{K}}$ copies of this

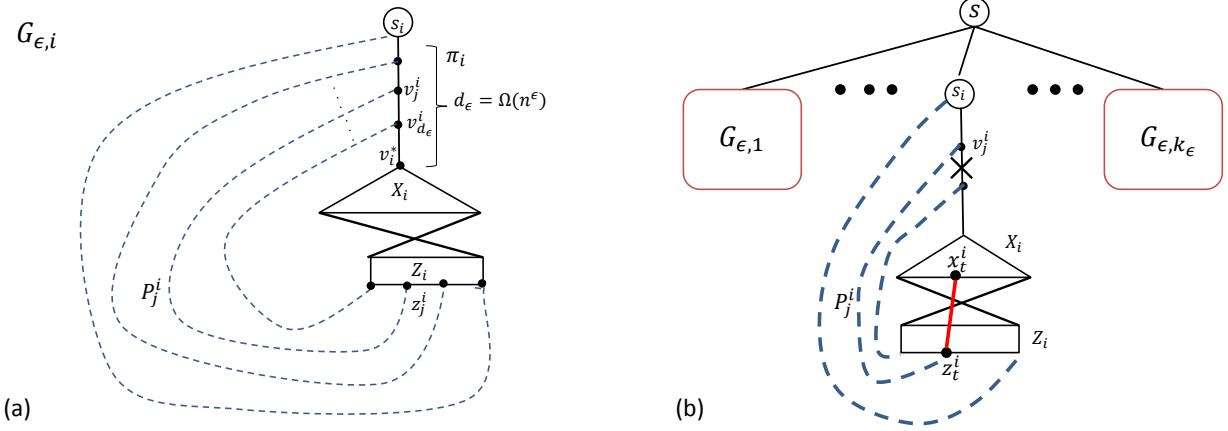


Figure 10: Schematic description of the lower bound construction for ϵ FT-BFS. (a) The graph $G_{\epsilon,i}$. The dashed lines corresponds to the collection monotone increasing lengths of paths P_j^i . The set of vertices X_i contains $\Omega(n^{2\epsilon})$ vertices that is fully connected to the collection of $\Omega(n^\epsilon)$ vertices Z_i . (b) The graph $G = (V, E)$ consists of $k_\epsilon = \Omega(n^{1-2\epsilon})$ copies of the graph $G_{\epsilon,i}$ connected to the source vertex s . The bold dashed red edge (x_t^i, z_j^i) is required upon failure of the edge e_j^i .

subgraph: $k_{\epsilon,\mathcal{K}}$ copies per source $s \in S$. Consider now the (i,j) 'th copy $G_{\epsilon,\mathcal{K}}^{i,j}$ for $i \in \{1, \dots, \mathcal{K}\}$ and $j \in \{1, \dots, k_{\epsilon,\mathcal{K}}\}$ (all copies are identical). This subgraph consists of three main components. The first is a path $\pi_{i,j} = [s_{i,j} = v_1^{i,j}, \dots, v_{d_{\epsilon,\mathcal{K}}}^{i,j} = v_{i,j}^{*,i}]$ of length $d_{\epsilon,\mathcal{K}}$. The second component consists of a node set $Z_{i,j} = \{z_1^{i,j}, \dots, z_{d_{\epsilon,\mathcal{K}}}^{i,j}\}$ and a collection of $d_{\epsilon,\mathcal{K}}$ disjoint paths of decreasing length, $P_1^{i,j}, \dots, P_{d_{\epsilon,\mathcal{K}}}^{i,j}$, where $P_\ell^{i,j} = [v_\ell^{i,j} = p_{\ell,1}^{i,j}, \dots, p_{\ell,t_\ell}^{i,j} = z_\ell^{i,j}]$ connects $v_\ell^{i,j}$ on $\pi_{i,j}$ with $z_\ell^{i,j}$ and its length is $t_\ell = |P_\ell^{i,j}| = 6 + 2(d_{\epsilon,\mathcal{K}} - \ell)$, for every $\ell \in \{1, \dots, d_{\epsilon,\mathcal{K}}\}$.

Altogether, the set of nodes in these paths, $Q_{i,j} = \bigcup_{\ell=1}^{d_{\epsilon,\mathcal{K}}} V(P_\ell^{i,j})$, is of size $|Q_{i,j}| = d_{\epsilon,\mathcal{K}}^2 + 5d_{\epsilon,\mathcal{K}} = \Theta((n/\mathcal{K})^{2\epsilon})$. This completes the description of the $G_{\epsilon,\mathcal{K}}^{i,j}$ subgraph.

For every $j \in \{1, \dots, k_{\epsilon,\mathcal{K}}\}$, let X_j be a set of $\Omega(n/k_{\epsilon,\mathcal{K}})$ vertices. The X_j vertices are connected via a star to a vertex \tilde{v}_j . The latter is connected to the terminal node $v_{i,j}^{*,i} \in \pi_{i,j}$ for every $i \in \{1, \dots, \mathcal{K}\}$. Formally, these edges are defined by $\widehat{E}(X_j) = \{(\tilde{v}_j, x) \mid x \in X_j\} \cup \{(\tilde{v}_j, v_{i,j}^{*,i}) \mid i \in \{1, \dots, \mathcal{K}\}\}$. In addition, the X_j vertices are fully connected to each of the vertices in $Z_j = \bigcup_{i=1}^{\mathcal{K}} Z_{i,j}$. Formally, G contains $k_{\epsilon,\mathcal{K}}$ complete bipartite graphs $B_j = B(X_j, Z_j)$ for every $j \in \{1, \dots, k_{\epsilon,\mathcal{K}}\}$. Finally, every source vertex $s_i \in S$ is connected to $s_{i,j}$ for every $j \in \{1, \dots, k_{\epsilon,\mathcal{K}}\}$. Overall, the vertices of $G = (V, E)$ are given by

$$V(G) = \bigcup_{i=1}^{\mathcal{K}} \bigcup_{j=1}^{k_{\epsilon,\mathcal{K}}} V(G_{\epsilon,\mathcal{K}}^{i,j}) \cup \bigcup_{j=1}^{k_{\epsilon,\mathcal{K}}} (X_j \cup \{\tilde{v}_j\}) \cup S$$

and the edges are

$$E(G) = \bigcup_{i=1}^{\mathcal{K}} \bigcup_{j=1}^{k_{\epsilon, \mathcal{K}}} (E(G_{\epsilon, \mathcal{K}}^{i,j}) \cup \{(s_i, s_{i,j})\}) \cup \bigcup_{j=1}^{k_{\epsilon, \mathcal{K}}} (E(B_j) \cup \widehat{E}(X_j)).$$

Observation 5.5. (a) $|V(G_{\epsilon, \mathcal{K}}^{i,j})| = \Theta((n/\mathcal{K})^{2\epsilon})$.

(b) The constants can be set precisely, so that $|V(G)| = n$ and $|X_j| = \Theta(n/k_{\epsilon, \mathcal{K}})$.

(c) $|E(G)| \geq |EB| = \Omega(\mathcal{K}^{1-\epsilon} \cdot n^{1+\epsilon})$.

Let $\Pi = \bigcup_{i=1}^{\mathcal{K}} \bigcup_{j=1}^{k_{\epsilon, \sigma}} E(\pi_{i,j})$. Since each $\pi_{i,j}$ is of length $d_{\epsilon, \sigma}$, overall

$$|E(\Pi)| \geq \lfloor \mathcal{K}^\epsilon \cdot n^{1-\epsilon} / 5 \rfloor. \quad (14)$$

Recall that $e_\ell^{i,j} = (v_\ell^{i,j}, v_{\ell+1}^{i,j})$ is the ℓ 'th edge on $\pi_{i,j}$.

The reasoning goes in a very similar way to the single source case. In particular, we show that the edges of Π are costly in the sense that unless they are reinforced in H , they require the introduction of many edges into H .

For ease of analysis, the edges of the bipartite graphs $B = \bigcup_{j=1}^{k_{\epsilon, \mathcal{K}}} B_j$ are partitioned into $|\Pi|$ disjoint sets $E_\ell^{i,j}$, each corresponds to one particular edge $e_\ell^{i,j} \in \Pi$. Let $E_\ell^{i,j} = \{(x_t^j, z_\ell^{i,j}) \mid x_t^j \in X_j\}$, hence $|E_\ell^{i,j}| = |X_j| = \Omega(n/k_{\epsilon, \mathcal{K}})$ and $EB = \bigcup_{i=1}^{\mathcal{K}} \bigcup_{j=1}^{k_{\epsilon, \sigma}} \bigcup_{\ell=1}^{d_{\epsilon, \sigma}} E_\ell^{i,j}$. For every ϵ FT-MBFS structure $H \subseteq G$, let $E'(H)$ be the set of at most $\lfloor \mathcal{K}^\epsilon \cdot n^{1-\epsilon} / 6 \rfloor = O(\mathcal{K}^\epsilon \cdot n^{1-\epsilon})$ reinforced edges in H (which is the maximum allowed number of reinforced edges by the statement of the theorem).

Claim 5.6. For every $e_\ell^{i,j} \in \Pi \setminus E'(H)$, it holds that $E_\ell^{i,j} \subseteq H$.

Proof: Assume, towards contradiction, that H does not contain $e' = (x_t^j, z_\ell^{i,j}) \in E_\ell^{i,j}$ for some $x_t^j \in X_j$. Note that upon the failure of the edge $e_\ell^{i,j} = (v_\ell^{i,j}, v_{\ell+1}^{i,j}) \in \pi_{i,j}$, the unique $s_i - x_t^j$ shortest path connecting s_i and x_t^j in $G \setminus \{e_\ell^{i,j}\}$ is $P'_{i,j,\ell} = \pi[s_i, v_\ell^{i,j}] \circ P_\ell^{i,j} \circ e'$, and all other alternatives are strictly longer. Note that the paths from the other sources $s_{i'} \in S \setminus \{s_i\}$ are not helpful as well. Since $e' \notin H$, also $P'_{i,j,\ell} \not\subseteq H$, and therefore $\text{dist}(s_i, x_t^j, G \setminus \{e_\ell^{i,j}\}) < \text{dist}(s_i, x_t^j, H \setminus \{e_\ell^{i,j}\})$, in contradiction to the fact that H is an ϵ FT-MBFS structure. It follows $E_\ell^{i,j} \subseteq H$. ■

We are now ready to complete of Theorem 5.4. Consider any ϵ FT-MBFS structure H . By Eq. (14) and by the bound on the number of reinforced edges as stated in Thm. 5.4, even if all reinforced edges $E'(H)$ are taken from the edges of Π (i.e., $E'(H) \subseteq E(\Pi)$), there are still $|E(\Pi) \setminus E'(H)| = \Omega(\mathcal{K}^\epsilon \cdot n^{1-\epsilon})$ edges in Π that remain unreinforced. By Cl. 5.6, each such edge $e_\ell^{i,j}$ requires that a subset of $\Omega(n/k_{\epsilon, \mathcal{K}})$ edges $E_\ell^{i,j}$ to be included in H . It then holds that $|E(H)| \geq \Omega(n/k_{\epsilon, \mathcal{K}}) \cdot |\Pi \setminus E'(H)| = \Omega(\mathcal{K}^{1-\epsilon} \cdot n^{1+\epsilon})$ as required. The theorem follows. ■

References

- [1] B. Awerbuch and Y. Azar. Buy-at-bulk network design. In *FOCS*, 1997.
- [2] S. Baswana and N. Khanna. Approximate shortest paths avoiding a failed vertex: near optimal data structures for undirected unweighted graphs. In *Algorithmica*, 66(1):18–50, 2013.
- [3] A. Bernstein and D. Karger. A nearly optimal oracle for avoiding failed vertices and edges. In *STOC*, 101–110, 2009.
- [4] G. Braunschvig, S. Chechik and D. Peleg. Fault tolerant additive spanners. In *WG*, 206–214, 2012.
- [5] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. Fault-tolerant spanners for general graphs. In *STOC*, 435–444, 2009.
- [6] C. Chekuri, M. T. Hajiaghayi, G. Kortsarz, M.R. Salavatipour. Approximation algorithms for nonuniform buy-at-bulk network design. In *SIAM J. Computing*, 39(5), 1772–1798, 2010.
- [7] A. Czumaj and H. Zhao. Fault-tolerant geometric spanners. *Discrete & Computational Geometry* **32**, (2003).
- [8] M. Dinitz and R. Krauthgamer. Fault-tolerant spanners: better and simpler. In *PODC*, 2011, 169–178.
- [9] F. Grandoni and V.V Williams. Improved Distance Sensitivity Oracles via Fast Single-Source Replacement Paths. In *FOCS*, 2012.
- [10] F. Grandoni and G.F. Italiano. Improved approximation for single-sink buy-at-bulk. In *Algorithms and Computation*: 111–120, 2006.
- [11] A. Gupta and A. Kumar and T. Roughgarden. Approximation via cost sharing: Simpler and better approximation algorithms for network design. In *Journal of the ACM (JACM)*, 54.3 (2007): 11.
- [12] C. Levkopoulos, G. Narasimhan, and M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *STOC*, 186–195, 1998.
- [13] T. Lukovszki. New results of fault tolerant geometric spanners. In *WADS*, 193–204, 1999.
- [14] M. Parter and D. Peleg. Sparse Fault-Tolerant BFS Trees. In *ESA*, 2013.
- [15] M. Parter and D. Peleg. Sparse Fault-Tolerant BFS Trees. In arxiv.org/pdf/1302.5401v1.pdf.
- [16] M. Parter. Vertex Fault Tolerant Additive Spanners. In *DISC*, 2014.
- [17] L. Roditty and U. Zwick. Replacement paths and k simple shortest paths in unweighted directed graphs. *ACM Trans. Algorithms*, 2012.
- [18] F.S. Salman, J. Cheriyan, R. Ravi, S. Subramanian. Buy-at-bulk network design: Approximating the single-sink edge installation problem. *SODA*, 1997.

- [19] D.D Sleator and R.E. Tarjan. A data structure for dynamic trees. *STOC*, 1981.
- [20] M. Thorup and U. Zwick. Approximate distance oracles. *J. ACM*, 52:1–24, 2005.
- [21] O. Weimann and R. Yuster. Replacement paths via fast matrix multiplication. *FOCS*, 2010.